

EXPRESS MAIL MAILING LABEL NUMBER EY 3/8623 698 US

PATENT

10830.0095.NPUS00

APPLICATION FOR UNITED STATES LETTERS PATENT

for

REPLICATION OF SNAPSHOT USING A FILE SYSTEM COPY
DIFFERENTIAL

By

Philippe Armangau
Milena Bergant

BACKGROUND OF THE INVENTION

1. Limited Copyright Waiver

A portion of the disclosure of this patent document contains computer code listings and command formats to which the claim of copyright protection is made. The copyright owner has no objection to the facsimile reproduction by any person of the patent document or the patent disclosure, as it appears in the U.S. Patent and Trademark Office patent file or records, but reserves all other rights whatsoever.

2. Field of the Invention

The present invention relates generally to data storage systems, and more particularly to a snapshot copy facility for a data storage system.

3. Description of the Related Art

Snapshot copies of a file system have been used for a variety of data processing and storage management functions such as storage backup, transaction processing, and software debugging. A snapshot copy facility, for example, stores a plurality of snapshot copies of a production file system. Each of the snapshot copies is a prior state of the production file system at a respective point in time. The snapshot copy facility has a bit map, a block map, and a save volume for each snapshot copy. The bit map for each snapshot copy indicates blocks of data in the production file system that have changed between the snapshot copy and a next snapshot copy of the production file system. The save volume for each snapshot copy stores the "before images" of the blocks that have changed between the snapshot copy and the next snapshot copy of the production file system. The block map for each snapshot copy provides the save volume block address given the block address of a block that has changed between the snapshot copy and the

1 next snapshot of the production file system.

3 SUMMARY OF THE INVENTION

4 In accordance with a first aspect of the invention, there is provided a method of
5 operating a snapshot copy facility that stores a plurality of snapshot copies of a
6 production file system. Each of the snapshot copies is a prior state of the production file
7 system at a respective point in time. The snapshot copy facility receives a request for the
8 difference between a specified older one of the snapshot copies and a specified younger
9 one of the snapshot copies. The snapshot copy facility responds to the request by
10 returning the difference between the specified older one of the snapshot copies and the
11 specified younger one of the snapshot copies.

12 In accordance with another aspect, the invention provides a method of operating a
13 snapshot copy facility that stores a plurality of snapshot copies of a production file
14 system. Each of the snapshot copies is a prior state of the production file system at a
15 respective point in time. The snapshot copy facility has an index for each snapshot copy
16 for indicating blocks of data in the production file system that have changed between the
17 snapshot copy and a next snapshot copy of the production file system. The method
18 includes scanning the indices for a sequence of the snapshot copies to determine the
19 blocks that have changed between an older one of the snapshot copies and a younger one
20 of the snapshot copies. The sequence of the snapshot copies includes the older one of the
21 snapshot copies and each of the snapshot copies that is both younger than the older one of
22 the snapshot copies and older than the younger one of the snapshot copies.

1 In accordance with another aspect, the invention provides a method of operating a
2 snapshot copy facility that stores a plurality of snapshot copies of a production file
3 system. Each of the snapshot copies is a prior state of the production file system at a
4 respective point in time. The snapshot copy facility has a first index for each snapshot
5 copy for indicating blocks of data in the production file system that have changed
6 between the snapshot copy and a next snapshot copy of the production file system and
7 that have a "before image" saved for said each snapshot copy. The snapshot copy facility
8 has a second index for each snapshot copy for indicating blocks of data that are not in use
9 in the snapshot copy. The method includes responding to a request for the difference
10 between a specified older one of the snapshot copies and a specified younger one of the
11 snapshot copies by accessing the second index for the specified younger one of the
12 snapshot copies to determine blocks of data in the production file system that are in use in
13 the specified younger one of the snapshot copies, and for blocks of data in the production
14 file system that are in use in the specified younger one of the snapshot copies, accessing
15 at least one of the first indices for a sequence of the snapshot copies to determine blocks
16 that have changed between an older one of the snapshot copies and a younger one of the
17 snapshot copies. The sequence of the snapshot copies includes the older one of the
18 snapshot copies and each of the snapshot copies that is both younger than the older one of
19 the snapshot copies and older than the younger one of the snapshot copies.

20 In accordance with yet another aspect, the invention provides a method of
21 operating a network file server. The network file server has a snapshot copy facility for
22 storing a plurality of snapshot copies of a production file system. Each of the snapshot
23 copies is a prior state of the production file system at a respective point in time. The

1 network file server receives a request for an update to a specified snapshot copy of the
2 production file system. The network file server responds to the request by checking
3 whether the snapshot copy facility contains the specified snapshot copy of the production
4 file system, and upon finding that the snapshot copy facility contains the specified
5 snapshot copy of the production file system, the network file server returns the difference
6 between the specified snapshot copy of the production file system and a more recent
7 snapshot copy of the production file system.

8 In accordance with yet another aspect, the invention provides a method of
9 operation in a data processing network having a client and a network file server. The
10 network file server stores a plurality of snapshot copies of a production file system. Each
11 of the snapshot copies is a prior state of the production file system at a respective point in
12 time. The client has a local version of an older one of the snapshot copies. The method
13 provides the client with a younger one of the snapshot copies. The method includes the
14 network file server determining the difference between the younger one of the snapshot
15 copies and the older one of the snapshot copies. The network file server transmits the
16 difference between the younger one of the snapshot copies and the older one of the
17 snapshot copies to the local version of the older one of the snapshot copies. The
18 difference between the younger one of the snapshot copies and the older one of the
19 snapshot copies is written into the local version of the older one of the snapshot copies to
20 produce a local version of the younger one of the snapshot copies.

21 In accordance with another aspect, the invention provides a snapshot copy
22 facility. The snapshot copy facility includes storage for storing a plurality of snapshot
23 copies of a production file system. Each of the snapshot copies is a prior state of the

1 production file system at a respective point in time. The snapshot copy facility also
2 includes at least one processor programmed for receiving a request for the difference
3 between a specified older one of the snapshot copies and a specified younger one of the
4 snapshot copies; and for responding to the request by returning the difference between
5 the specified older one of the snapshot copies and the specified younger one of the
6 snapshot copies.

7 In accordance with yet another aspect, the invention provides a snapshot copy
8 facility. The snapshot copy facility includes storage for storing a plurality of snapshot
9 copies of a production file system. Each of the snapshot copies is a prior state of the
10 production file system at a respective point in time. The snapshot copy facility includes
11 an index for each snapshot copy for indicating blocks of data in the production file
12 system that have changed between the snapshot copy and a next snapshot copy of the
13 production file system. The snapshot copy facility also includes at least one processor
14 programmed for scanning the indices for a sequence of the snapshot copies to determine
15 the blocks that have changed between an older one of the snapshot copies and a younger
16 one of the snapshot copies, the sequence of the snapshot copies including the older one of
17 the snapshot copies and each of the snapshot copies that is both younger than the older
18 one of the snapshot copies and older than the younger one of the snapshot copies.

19 In accordance with another aspect, the invention provides a snapshot copy
20 facility. The snapshot copy facility includes storage for storing a plurality of snapshot
21 copies of a production file system. Each of the snapshot copies is a prior state of the
22 production file system at a respective point in time. The snapshot copy facility has a first
23 index for each snapshot copy for indicating blocks of data in the production file system

1 that have changed between the snapshot copy and a next snapshot copy of the production
2 file system and that have a "before image" for the snapshot copy stored in the storage.
3 The snapshot copy facility has a second index for each snapshot copy for indicating
4 blocks of data that are not in use in the snapshot copy. The snapshot copy facility also
5 has at least one processor programmed for responding to a request for the difference
6 between a specified older one of the snapshot copies and a specified younger one of the
7 snapshot copies by accessing the second index for the specified younger one of the
8 snapshot copies to determine blocks of data in the production file system that are in use in
9 the specified younger one of the snapshot copies, and for blocks of data in the production
10 file system that are in use in the specified younger one of the snapshot copies, accessing
11 at least one of the first indices for a sequence of the snapshot copies to determine blocks
12 that have changed between an older one of the snapshot copies and a younger one of the
13 snapshot copies. The sequence of the snapshot copies includes the older one of the
14 snapshot copies and each of the snapshot copies that is both younger than the older one of
15 the snapshot copies and older than the younger one of the snapshot copies.

16 In accordance with still another aspect, the invention provides a network file
17 server including a snapshot copy facility for storing a plurality of snapshot copies of a
18 production file system. Each of the snapshot copies is a prior state of the production file
19 system at a respective point in time. The network file server is programmed for receiving
20 a request for an update to a specified snapshot copy of the production file system, and
21 responding to the request by checking whether the snapshot copy facility contains the
22 specified snapshot copy of the production file system, and upon finding that the snapshot
23 copy facility contains the specified snapshot copy of the production file system, returning

1 the difference between the specified snapshot copy of the production file system and a
2 more recent snapshot copy of the production file system.

3 In accordance with another aspect, the invention provides a program storage
4 device containing a program for a snapshot copy facility. The snapshot copy facility
5 stores a plurality of snapshot copies of a production file system. Each of the snapshot
6 copies is a prior state of the production file system at a respective point in time. The
7 program is executable for responding to a request for the difference between a specified
8 older one of the snapshot copies and a specified younger one of the snapshot copies by
9 returning the difference between the specified older one of the snapshot copies and the
10 specified younger one of the snapshot copies.

11 In accordance with yet another aspect, the invention provides a program storage
12 device containing a program for a snapshot copy facility. The snapshot copy facility has
13 a plurality of snapshot copies of a production file system. Each of the snapshot copies is
14 a prior state of the production file system at a respective point in time. The snapshot
15 copy facility also has an index for each snapshot copy for indicating blocks of data in the
16 production file system that have changed between the snapshot copy and a next snapshot
17 copy of the production file system. The program is executable for scanning the indices
18 for a sequence of the snapshot copies to determine the blocks that have changed between
19 an older one of the snapshot copies and a younger one of the snapshot copies, the
20 sequence of the snapshot copies including the older one of the snapshot copies and each
21 of the snapshot copies that is both younger than the older one of the snapshot copies and
22 older than the younger one of the snapshot copies.

1 In accordance with another aspect, the invention provides a program storage
2 device containing a program for a snapshot copy facility. The snapshot copy facility has
3 a plurality of snapshot copies of a production file system. Each of the snapshot copies is
4 a prior state of the production file system at a respective point in time. The snapshot
5 copy facility has a first index for each snapshot copy for indicating blocks of data in the
6 production file system that have changed between the snapshot copy and a next snapshot
7 copy of the production file system and that have a "before image" for the snapshot copy
8 stored in the snapshot copy facility. The snapshot copy facility has a second index for
9 each snapshot copy for indicating blocks of data that are not in use in the snapshot copy.
10 The program is executable for responding to a request for the difference between a
11 specified older one of the snapshot copies and a specified younger one of the snapshot
12 copies by accessing the second index for the specified younger one of the snapshot copies
13 to determine blocks of data in the production file system that are in use in the specified
14 younger one of the snapshot copies, and for blocks of data in the production file system
15 that are in use in the specified younger one of the snapshot copies, accessing at least one
16 of the first indices for a sequence of the snapshot copies to determine blocks that have
17 changed between an older one of the snapshot copies and a younger one of the snapshot
18 copies. The sequence of the snapshot copies includes the older one of the snapshot
19 copies and each of the snapshot copies that is both younger than the older one of the
20 snapshot copies and older than the younger one of the snapshot copies.

21 In accordance with a final aspect, the invention provides a program storage device
22 containing a program for a network file server. The network file server includes a
23 snapshot copy facility for storing a plurality of snapshot copies of a production file

1 system. Each of the snapshot copies is a prior state of the production file system at a
2 respective point in time. The program is executable for receiving a request for an update
3 to a specified snapshot copy of the production file system, and responding to the request
4 by checking whether the snapshot copy facility contains the specified snapshot copy of
5 the production file system, and upon finding that the snapshot copy facility contains the
6 specified snapshot copy of the production file system, returning the difference between
7 the specified snapshot copy of the production file system and a more recent snapshot
8 copy of the production file system.

9

10 **BRIEF DESCRIPTION OF THE DRAWINGS**

11 Other objects and advantages of the invention will become apparent upon reading
12 the following detailed description with reference to the accompanying drawings wherein:

13 FIG. 1 is a block diagram of an IP network including multiple hosts and multiple
14 data mover computers;

15 FIG. 2 is a block diagram showing a primary data mover distributing remote copy
16 data to multiple secondary data movers in the IP network by establishing a Transmission
17 Control Protocol (TCP) connection with each of the secondary data movers;

18 FIG. 3 is a block diagram showing a primary data mover distributing remote copy
19 data to multiple data movers through forwarder data movers;

20 FIG. 4 is a block diagram showing a shared save volume used to buffer local copy
21 data transmitted from a primary data mover to a secondary data mover;

22 FIG. 5 is a block diagram showing a primary save volume and a secondary save
23 volume;

FIG. 6 is a flowchart showing local replication in the system of FIG. 4;

FIG. 7 is a flowchart showing remote replication in the system of FIG. 5;

FIG. 8 is a block diagram of a primary site, including layered programming in a primary data mover;

FIG. 9 is a block diagram of a secondary site, including layered programming in a secondary data mover;

FIG. 10 is a flowchart of a process of replication at the primary site of FIG. 8;

FIG. 11 is a flowchart of a procedure for producing a new remote copy of a primary file system concurrent with ongoing replication and multicasting of modifications to the primary file system;

FIG. 12 is a flowchart of an IP-replication send-thread introduced in FIG. 8;

FIG. 13 is a block diagram of a volume multicast level in the data mover programming of FIG. 8 and FIG. 9;

FIG. 14 is a block diagram of the RCP level in the primary data mover programming of FIG. 8;

FIG. 15 is a block diagram of the RCP level in the secondary data mover programming of FIG. 9;

FIG. 16 is a block diagram of a RPC forwarder at the RPC level in a forwarder data mover;

FIG. 17 is a flowchart of an inbound RCP session in the secondary data mover;

FIG. 18 is a block diagram showing a forwarder data mover performing local replication;

1 FIG. 19 is a block diagram showing the sharing of a data mover's single TCP port
2 for RCP connections with Hypertext Transfer Protocol (HTTP) connections;

3 FIG. 20 is a block diagram showing the state of a remote replication system
4 before a disaster at the primary site;

5 FIG. 21 is a flowchart of a failover and recovery method performed upon the
6 remote replication system of FIG. 20 in order to recover from the disaster at the primary
7 site;

8 FIG. 22 is a block diagram showing the state of the remote replication system of
9 FIG. 20 after failover from the primary site to the secondary site;

10 FIG. 23 is a flowchart of a subroutine of failover with a checkpoint and without
11 sync, as used in the flowchart of FIG. 21;

12 FIG. 24 shows a block map that can be used to create a snapshot copy of the
13 secondary file system at a restart point during the failover of FIG. 23;

14 FIG. 25 shows a snapshot copy facility that can be used to create the snapshot
15 copy of the secondary file system at the restart point during the failover of FIG. 23;

16 FIG. 26 is a flowchart of a procedure for writing a block of data to a production
17 file system in the snapshot copy facility of FIG. 25;

18 FIG. 27 is a flowchart of a procedure for reading a block of data from a snapshot
19 file system in the snapshot copy facility of FIG. 25;

20 FIG. 28 is a flowchart of a subroutine to resync the primary file system with the
21 secondary file system, as used in the flowchart of FIG. 21;

22 FIG. 29 shows the state of the remote replication system of FIG. 20 during the
23 resync procedure;

1 FIG. 30 is a flowchart of a subroutine to failback to the primary file system, as
2 used in the flowchart of FIG. 21;

3 FIG. 31 is a flowchart of execution of a failover command having a sync option
4 and a checkpoint option;

5 FIG. 32 is a flowchart of a subroutine for failover without sync, as used in the
6 flowchart of FIG. 31;

7 FIG. 33 is a subroutine for failover with sync, as used in the flowchart of FIG. 31;

8 FIGS. 34 and 35 comprise a flowchart of a procedure for copying snapshot delta
9 for snapshots L to M;

10 FIGS. 36 shows a block diagram of a data network in which snapshot deltas are
11 transmitted over a wide-area network from a network file server to a local file server in
12 order to update the local file system as needed;

13 FIGS. 37 and 38 comprise a flowchart of a procedure for replicating snapshots in
14 the data network of FIG. 36;

15 FIG. 39 is a modified version of the flowchart of FIG. 34, showing how to copy a
16 snapshot delta from a snapshot copy facility that keeps track of blocks of a production
17 file system that are not in use in the snapshot copies of the production file system;

18 FIG. 40 shows a specific instance of the flowchart of FIG. 39 for a snapshot copy
19 facility that uses respective meta bit maps for indicating the blocks in a production file
20 system that are not used in the snapshot copies of the production file system;

21 FIG. 41 is a flowchart for writing a specified data block to the production file
22 system for a snapshot copy facility that uses a meta bit map for indicating the blocks in a
23 production file system that are not presently in use;

Remote Data Replication Facility

FIG. 1 shows an IP network 220 including multiple network file servers 221, 222, and multiple hosts 223, 224, 225. The hosts and network file servers, for example, can be distributed world wide and linked via the Internet. Each of the network file servers 221, 222, for example, has multiple data movers 226, 227, 228, 232, 233, 234, for moving data between the IP network 220 and the cached disk array 229, 235, and a control station 230, 236 connected via a dedicated dual-redundant data link 231, 237 among the data movers for configuring the data movers and the cached disk array 229, 235. Further details regarding the network file servers 221, 222 are found in Vahalia et al., U.S. Patent 5,893,140, incorporated herein by reference.

In operation, it is desired for each of the network file servers 221, 222 to provide read-only access to a copy of the same file system. For example, each of the network file servers could be programmed to respond to user requests to access the same Internet site. The IP network 220 routes user requests to the network file servers 221, 222 in the same continent or geographic region as the user. In this fashion, the user load is shared among the network file servers.

In the wide-area network of FIG. 1, it is desired to perform read-write updating of the respective file system copies in the network file servers 221, 222 while permitting concurrent read-only access by the hosts. It is also desired to distribute the updates over the IP network.

There are a number of ways that updates could be distributed over the IP network from a primary data mover to multiple secondary data movers. As shown in FIG. 2, for example, a primary data mover establishes a connection 242, 243, 244 in accordance with

1 the industry-standard Transmission Control Protocol (TCP) over the IP network 220 to
2 each secondary data mover 245, 246, 247, and then concurrently sends the updates to
3 each secondary data mover over the TCP connection. When the updates need to be
4 distributed to a large number of secondary data movers, however, the amount of time for
5 distributing the updates may become excessive due to limited resources (CPU execution
6 cycles, connection state, or bandwidth) of the primary data mover 241. One way of
7 extending these limited resources would be to use existing IP routers and switches to
8 implement "fan out" from the primary data mover 241 to the secondary data movers 245,
9 246, 247. Still, a mechanism for reliability should be layered over the Internet Protocol.

10 FIG. 3 shows that the time for distributing updates from a primary data mover 251
11 to a large number of secondary data movers 254, 255, 256, 257 can be reduced by using
12 intermediate data movers 252, 253 as forwarders. The primary data mover 251 sends the
13 updates to the forwarder data movers 252, 253, and each of the forwarder data movers
14 sends the updates to a respective number of secondary data movers. The forwarder data
15 movers 252, 253 may themselves be secondary data movers; in other words, each may
16 apply the updates to its own copy of the replicated read-only file system. The distribution
17 from the primary data mover 251 to the forwarder data movers 252, 253 can be done in a
18 fashion suitable for wide-area distribution (such as over TCP connections). The
19 forwarding method of replication of FIG. 3 also has the advantage that the distribution
20 from each forwarder data mover to its respective data movers can be done in a different
21 way most suitable for a local area or region of the network. For example, some of the
22 forwarder data movers could use TCP connections, and others could use a combination of

1 TCP connections for control and UDP for data transmission, and still other forwarders
2 could be connected to their secondary data movers by a dedicated local area network.

3 For implementing the replication method of FIG. 3 over the Internet Protocol,
4 there are a number of desired attributes. It is desired to maintain independence between
5 the primary data mover and each of the secondary data movers. For example, a new
6 secondary data mover can be added at any time to replicate an additional remote copy.
7 The primary data mover should continue to function even if a secondary data mover
8 becomes inoperative. It is also desired to maintain independence between the replication
9 method and the IP transport method. Replication should continue to run even if the IP
10 transport is temporarily inactive. It is desired to recover in a consistent fashion from a
11 panic or shutdown and reboot. A record or log of the progress of the replication can be
12 stored for recovery after an interruption. It is desired to build re-usable program blocks
13 for the replication function, so that the program blocks for the replication function can be
14 used independent of the location of the primary file system or its replicas.

15 In a preferred implementation, independence between the replication process, the
16 IP transport method, and the primary file system being replicated, is ensured by use of a
17 save volume. The save volume is a buffer between the data producer (i.e., the host or
18 application updating the primary file system), the replication process, and the data
19 consumer (the secondary data movers). The save volume stores the progress of the
20 replication over the Internet Protocol so as to maintain the consistency of the replication
21 process upon panic, reboot, and recovery. The transport process need not depend on any
22 "in memory" replication information other than the information in the save volume, so as

1 to permit the replication process to be started or terminated easily on any data mover for
2 load shifting or load balancing.

3 When a save volume is used, it can be shared between a primary data mover and a
4 secondary data mover in the case of local file system replication, or a primary copy of the
5 shared volume can be kept at the primary site, and a secondary copy of the shared volume
6 can be kept at the secondary site, in the case of remote file system replication.

7 For the case of local file system replication, FIG. 4 shows a primary site including
8 a primary data mover 260 managing access to a primary file system 261, and a secondary
9 data mover 262 managing access to a secondary file system 263 maintained as a read-
10 only copy of the primary file system 261. A save volume 264 is shared between the
11 primary data mover 260 and the secondary data mover 262. This sharing is practical
12 when the secondary site is relatively close to the primary site. A redo log 265 records a
13 log of modifications to the primary file system 261 during the replication process for
14 additional protection from an interruption that would require a reboot and recovery.

15 Local replication can be used to replicate files within the same network file
16 server. For example, in the network file server 221 in FIG. 1, the primary data mover
17 could be the data mover 226, the secondary data mover could be the data mover 227, the
18 save volume could be stored in the cached disk array 229, and replication control
19 messages could be transmitted between the data movers over the data link 231.

20 For the case of remote file system replication, FIG. 5 shows a primary site
21 including a primary data mover 270 managing access to a primary file system 271, and a
22 secondary data mover 272 managing access to a secondary file system 273 maintained as
23 a read-only copy of the primary file system 271. The primary site includes a primary

1 save volume 274, and the remote site includes a secondary save volume 275. A redo log
2 276 records a log of modifications to the primary file system 271 during the replication
3 process for additional protection from an interruption that would require a reboot and
4 recovery.

5 FIG. 6 shows a method of operating the system of FIG. 4 for local replication. In
6 a first step 281, the primary data mover migrates a copy of the primary file system to
7 create a secondary file system at the secondary site in such a way to permit concurrent
8 write access to the primary file system. The migration, for example, may use the method
9 shown in FIG. 17 of Ofek U.S. Patent 5,901,327, in which a bit map indicates remote
10 write pending blocks. Alternatively, the migration may use a snapshot copy mechanism,
11 for example, as described in Kedem, U.S. Patent 6,076,148, in which a bit map indicates
12 the blocks that have changed since the time of snap-shotting of the primary file system.
13 The snapshot method is preferred, because it is most compatible with the delta set
14 technique for remote copy of subsequent modifications. For example, a snapshot
15 manager creates a snapshot copy of the primary file system, as will be further described
16 below with reference to FIGS. 25 to 27. The migration may involve taking a first
17 snapshot copy of the primary file system, copying the first snapshot copy to the
18 secondary file system, starting replication and creating a second snapshot copy of the
19 primary file system, and then copying the differential between the first snapshot copy and
20 the second snapshot copy from the primary file system to the secondary file system, in a
21 fashion similar to that described further below with reference to steps 536 and 537 in
22 FIG. 28. In any event, it is desired for the secondary file system to become a copy of the
23 state of the primary file system existing at some point of time, with any subsequent

1 modifications of the primary file system being transferred through the shared save
2 volume.

3 In step 282, the primary data mover writes subsequent modifications of the
4 primary file system to the shared save volume. In step 283, the secondary data mover
5 reads the subsequent modifications from the shared save volume and writes them to the
6 secondary file system. In step 284, the secondary data mover provides user read-only
7 access to consistent views of the secondary file system, by integrating the subsequent
8 revisions into the secondary file system while providing concurrent read-only access to
9 the secondary file system. This can be done by using a remote data facility to maintain a
10 remote copy of a pair of delta volumes and to switch between the delta volumes as
11 described in Suchitra Raman, et al., U.S. Patent Application Ser. No. 10/147,751 filed
12 May 16, 2002, entitled "Replication of Remote Copy Data for Internet Protocol (IP)
13 transmission, incorporated herein by reference. This can also be done by successively
14 accumulating delta sets at the secondary, and accessing a block index or map of updated
15 data blocks in the delta set.

16 Each delta set, for example, is a set of consistent updates to blocks of the
17 secondary file system. The updated blocks are included in "delta chunks" transmitted
18 from the primary file server to the secondary file server. Each delta set includes a
19 number of delta chunks, and each delta chunk belongs to a single delta set. Each delta
20 chunk is transmitted with a header including a generation count, a delta set number, and a
21 chunk number within the delta set. When a delta set is accumulated, a block index or
22 map is created to identify the updated blocks in the delta set. Once the delta set has been
23 accumulated, a background integration process is started that writes the updates into the

1 secondary file system, and the progress of the background process is recorded by
2 updating the block index or map for the delta set when each updated block is written to
3 the secondary file system. During the background process, read access to the secondary
4 file system is permitted on a priority basis. For this read access to a specified block of
5 the file system, the secondary file server first accesses the block index or map of the delta
6 set, and if the specified block is in the delta set, the secondary file server returns the data
7 of the block from the delta set. If the specified block is not in the delta set, then the
8 secondary file server returns the data of the block from the secondary file system.

9 In FIG. 6, execution loops from step 284 back to step 282. In this fashion, the
10 secondary file system is updated from the primary site concurrently with read-only access
11 at the secondary site.

12 FIG. 7 shows a method of operating the system of FIG. 5 for remote replication.
13 In a first step 291, the primary data mover migrates a copy of the primary file system to
14 create a secondary file system at the secondary site, in a fashion similar to step 281 in
15 FIG. 6. In step 292, the primary data mover writes subsequent modifications of the
16 primary file system to the primary save volume, in a fashion similar to step 282 in FIG. 6.
17 In step 293, the modifications are copied from the primary save volume to the secondary
18 save volume, by transmitting delta chunks. In step 294, the secondary data mover reads
19 the modifications from the secondary save volume and writes them to the secondary file
20 system. In step 295, the secondary data mover provides user read-only access to
21 consistent views of the secondary file system, in a fashion similar to step 284 of FIG. 6.
22 Execution loops from step 295 back to step 292. In this fashion, the secondary file

1 system is remotely updated from the primary site concurrently with read-only access at
2 the secondary site.

3 FIG. 8 shows layered programming 300 for a primary data mover.. It is desired to
4 use layered programming in accordance with the International Standard Organization's
5 Open Systems Interconnection (ISO/OSI) model for networking protocols and distributed
6 applications. As is well known in the art, this OSI model defines seven network layers,
7 namely, the physical layer, the data link layer, the network layer, the transport layer, the
8 session layer, the presentation layer, and the application layer.

9 As shown in FIG. 8, the layered programming 300 includes a conventional
10 TCP/IP transport layer 301. The layers above the TCP/IP transport layer 301 include a
11 replication control protocol (RCP) session layer 302, a volume multicast presentation
12 layer 303, and an IP-FS (file system) copy send-thread 304 and an IP-replication send-
13 thread 305 at the program layer level. Over these program layers is a management and
14 configuration command interpreter (MAC_CMD) 306 for system operator set-up,
15 initiation, and supervisory control of the replication process.

16 In operation, the RCP layer 302 provides an application program interface (API)
17 for multicasting data over TCP/IP. RCP provides callback, acknowledgement (ACK),
18 and resumption of aborted transfers.

19 RCP provides the capability for a remote site to replicate and rebroadcast remote
20 copy data. The remote site functions as a router when it rebroadcasts the remote copy
21 data. RCP can also be used to replicate data locally within a group of data movers that
22 share a data storage system.

1 To create a new remote copy in response to a supervisory command, the
2 command interpreter 306 initiates execution of a replication module 310 if the replication
3 module is not presently in an active mode. Then, the command interpreter 306 invokes a
4 snapshot manager 308 to create a snapshot copy 309 of a primary file system volume
5 307. When the snapshot copy is created, the snapshot manager 308 obtains a current
6 delta set number from the replication module 310 and inserts the current delta set number
7 into the metadata of the snapshot. The current delta set number for the snapshot is all that
8 the secondary needs to identify modifications that are made subsequent to the creation of
9 the snapshot. In this fashion, any number of new remote copies can be created at various
10 times during operation of the replication module, with the snapshot process operating
11 concurrently and virtually independent of the replication module. For example,
12 whenever synchronization of a remote copy is lost, for example due to a prolonged
13 disruption of network traffic from the primary site to the remote site, a new remote copy
14 can be created to replace the unsynchronized remote copy.

15 Once the snapshot copy 309 is accessible, the command interpreter 306 initiates
16 execution of an instance of the IP-FS copy send-thread 304. The instance of the IP-FS
17 copy send-thread 304 reads data from the snapshot copy 309 and calls upon the volume
18 multicast layer 303 to multicast the remote copy data to all of the secondary data movers
19 where the remote copies are to be created. This can be a copy by extent, so there is no
20 copying of invalid or unused data blocks. For example, the volume multicast layer 303 is
21 given a copy command (@vol., length) specifying a volume and an extent to be copied,
22 and may also specify a group of destinations (an RCP group). The snapshot copy 309 of
23 the primary file system identifies the next valid block to be copied, and the number of

1 valid contiguous blocks following the next block. These blocks are copied at the logical
2 level, so it does not matter what physical structure is used for storing the secondary file
3 system at the secondary site. The copying is done locally, or by remote copy, for
4 example by transporting the data block over IP. The volume multicast layer 303 invokes
5 the RCP layer 302 to transport each data block.

6 During the remote copy process, whenever a modification is made to a block of
7 the primary file system volume 307, the replication module 310 logs an indication of the
8 modified block in a log 314 and later assembles the modification into a delta set chunk
9 written to a primary save volume 311. The replication module 310 logs the indications in
10 the log 314 on a priority or foreground basis as data is written to the primary file system
11 volume 307, and also logs boundaries between delta sets. The replication module 310
12 later reads the log 314 to read the indicated modifications from the primary file system
13 volume 307, assemble the indicated modifications into delta set chunks on a background
14 basis, and store the delta set chunks in a save volume chunk area of the save volume 311.
15 For example, the log is in the form of a queue of two bit-map tables, a new one of the
16 tables being written to coincident with write operations upon the primary file system
17 volume 307, and an old one of the tables being read to determine blocks to copy from the
18 primary file system to create a new delta set in the save volume 311. When the delta set
19 chunks become available for distribution from the save volume 311, the replication
20 module 310 updates the save volume mailbox area 312 by storing each delta set chunk
21 definition (@vol., length).

22 The IP-replication send-thread instance 305 polls the save volume mailbox area
23 312 to see if any delta set chunks have been stored in the save volume chunk area 313. If

1 so, then the thread instance calls upon the volume multicast layer 303 to multicast the
2 delta set chunks to the data movers that manage the storage of the respective remote file
3 system copies. For example, for each delta set chunk, the IP-replication send-thread
4 instance 305 issues a volume multicast command to the volume multicast layer 303.
5 When the chunk multicast is completed, the IP-replication send-thread instance 305
6 updates its context on the save volume 311 in the mailbox area 312. At reboot after an
7 interruption of multicast of a chunk, the IP-replication send-thread instance is able to
8 restart the multicast of the chunk. The IP-replication send-thread instance also is
9 responsible for retrying transmission of the chunk whenever the connection with the
10 secondary is interrupted.

11 FIG. 9 shows the layered programming 320 for a secondary data mover. The
12 programming includes a TCP/IP layer 321, an RCP layer 322, a volume multicast layer
13 323, and a management and configuration command interpreter (MAC_CMD) 324.
14 During creation of a new remote copy in a secondary file system volume 325, the volume
15 multicast layer 323 writes remote copy data from the primary data mover to the
16 secondary file system volume 325, and concurrently writes modifications (delta set
17 chunks) from the primary data mover to a save volume chunk area 326 of a secondary
18 save volume 327.

19 A header for the changes in a next version of the delta set is sent last, because
20 there is no guarantee of the order of receipt of the IP packets. The header of the delta set
21 includes a generation count, the number of delta blocks for the next version of the delta
22 set, a checksum for the header, and a checksum for the data of all the delta blocks. The
23 receiver checks whether all of the changes indicated in the header have been received.

1 Once a complete remote snapshot copy has been reconstructed in the secondary
2 file system volume 325, a playback module 328 is activated to read the modifications
3 from the save volume chunk area 326 and integrates them into the secondary file system
4 volume 325. From each delta-set chunk in the save volume area 326, the playback
5 module 328 gets the block address and number of contiguous blocks to be written to the
6 secondary file system volume. An access module 329 provides read-only access to a
7 consistent view of the secondary file system in the secondary file system volume 325.

8 FIG. 10 shows a procedure executed by the primary site of FIG. 8 to perform
9 replication of the primary file system. When replication is started in a first step 341, the
10 primary file system is paused to make it consistent. Migration of the primary file system
11 to the secondaries can then be started using a remote copy facility or snapshot manager.
12 Then, in step 342, concurrent write access to the primary file system is resumed, and all
13 modifications made on the primary file system are logged at the volume level on a
14 priority or foreground basis when each modification is made. In addition, a background
15 process of delta-set creation is initiated.

16 Two configurable triggers specify the rate of delta set creation: a timeout
17 parameter and a high water mark parameter. Whenever delta set creation is initiated, the
18 current time, as indicated by a real-time clock, is added to a configurable timeout interval
19 to produce the timeout parameter. The high water mark specifies an amount of modified
20 data, in megabytes. The first trigger that occurs will trigger the creation of a delta set.
21 The replication module creates the delta set by pausing the primary file system, copying
22 the modified blocks from the primary file system to the delta set volume, and then
23 resuming the primary file system. By logging indications of the modified blocks and

1 later copying the modified blocks, multiple modifications to the same block are
2 represented and transported once during a single delta set.

3 In step 343, the background process of delta set creation is temporarily suspended,
4 for example, by placing the process on a task queue that is periodically serviced. In step
5 344, execution of the delta set creation process is resumed. In step 345, the modification
6 size is compared to the high water mark. If the high water mark is not exceeded, then
7 execution continues to step 346. In step 346, the present value of the real-time clock is
8 compared to the timeout parameter. If the timeout parameter has not been exceeded, then
9 execution loops back to step 343. Otherwise, execution continues to step 347. Execution
10 also branches to step 347 from step 345 if the modification size is greater than the high
11 water mark.

12 In step 347, the primary file system is paused. In step 348, a new delta set is
13 created by starting the copying of modified blocks from the primary file system volume
14 to the new delta set. In step 349, the logging of new modifications into a new table is
15 started. In step 350, the time-out and high water mark is re-armed. In other words, a new
16 value for the timeout parameter is computed as the current real time plus the configurable
17 timeout interval, and the modification size is reset to indicate the size of the new
18 modifications. In step 351, the primary file system is resumed. Execution loops from
19 step 351 back to step 343 to suspend the background process of delta set creation.

20 To maintain the consistency of the delta set created in the primary save volume,
21 the primary file system could remain paused and not resumed in step 351 until the copy
22 process begun in step 348 is completed. Preferably, however, the copy process begun in
23 step 348 is a snapshot copy process, so that write access to the primary file system may

1 resume in step 351 before the copy process has been completed. For the example of the
2 modification log being a queue of two bit-map tables, when a write access to a block in
3 the primary file system is requested, the old bit map is accessed on a priority basis. If the
4 corresponding bit in the old bit map indicates a modified block in the primary file system
5 volume not yet copied to the save volume, then it is copied on a priority basis to the save
6 volume before the new write data is written to the primary file system volume. As soon
7 as a modified block has been copied from the primary file system volume to the save
8 volume, the corresponding bit in the old bit map is cleared. In this fashion, at the
9 completion of the copy process, the entire old table will be in a reset state, ready to be
10 used as the next new table.

11 When the copy process started in step 348 is completed, the replication module
12 sets the save volume mailbox area to show that a new delta set is ready for transmission.
13 Upon polling the mailbox area, the IP-replication send-thread finds that the new delta set
14 is ready for transmission, and invokes the volume multicast layer to transmit the delta set
15 to the secondary sites. After step 351, execution loops back to step 343.

16 FIG. 11 shows a flow chart of the overall procedure of creating a new remote
17 copy, either for the first time at a secondary site or as a replacement for a remote copy
18 that needs to be resynchronized with the primary file system. In a first step 352, the
19 snapshot manager creates a snapshot copy of the primary file system at the end of any
20 pending transaction upon the primary file system (e.g., when the primary file system
21 becomes consistent after it is paused in step 341 of FIG. 10 or in step 347 of FIG. 10.)
22 The replication module independently writes any subsequent modifications into a current
23 delta set for the next transaction.

1 In step 353, the snapshot manager obtains the current delta set number from the
2 replication module and inserts it into metadata of the snapshot copy. In step 354, the IP-
3 FS copy send-thread is started in order to send volume extents of the snapshot copy to the
4 secondary data mover, by invoking the volume multicast layer for each extent.

5 In step 355, when the IP-FS copy send-thread is finished, the primary data mover
6 sends a "start playback" signal to the secondary data mover. In step 356, the secondary
7 data mover receives the "start playback" signal from the primary data mover, and starts
8 the playback module. In step 357, playback module begins playback from the delta set
9 indicated by the delta set number in the snapshot metadata.

10 The playback module (328 in FIG. 23) at the secondary site integrates the delta
11 set modifications into secondary file system. Each time that a new delta set appears in
12 the secondary save volume, the modifications can be integrated into the secondary file
13 system, for example, by pausing the secondary file system, copying the modifications
14 from the secondary save volume into the secondary file system, and resuming the
15 secondary file system. Alternatively, a timeout interval and a high water mark value can
16 be configured for the secondary site, so that the modifications may be integrated into the
17 secondary file system at a rate less frequent than the rate at which the new delta sets
18 appear in the secondary save volume. In this case, the modifications from the secondary
19 save volume would not be integrated into the secondary file system until the timeout time
20 is reached unless the amount of modifications in the save volume reaches the high water
21 mark. As described above, the integration of the modifications can be performed
22 concurrently with read-only access to a consistent view of the secondary file system.

FIG. 12 shows a flowchart of the IP-replication send-thread (305 in FIG. 8). In a first step 361, the thread polls the primary save volume mailbox area. If the mailbox area indicates that there is not a new delta set chunk in the primary save volume area, then the thread is finished for the present task invocation interval. Execution of the thread is suspended in step 363, and resumed in step 364 at the next task invocation interval.

If the mailbox area indicates that there is a new delta set chunk in the primary save volume, then execution continues from step 362 to step 365. In step 365, the IP-replication send-thread issues a volume multicast command to broadcast or forward the delta set chunk to specified destination data movers. In step 366, if the multicast has been successful, then execution branches to step 367. In step 367, the IP-replication send-thread updates the primary save volume mailbox to indicate completion of the multicast, and execution continues to step 363 to suspend execution of the thread until the next task invocation interval.

In step 366, if the multicast is not successful, then execution continues to step 368 to test whether more than a certain number (N) of retries have been attempted. If not, then execution loops back to step 365 to retry the multicast of step 365. If more than N retries have been attempted, then execution continues from step 368 to step 369. In step 369, the IP-replication send-thread logs the error, and then in step 370, passes execution to an error handler.

FIG. 13 shows various objects defined by the volume multicast layer. The volume multicast layer provides multicast service to instances of a VolMCast object 370 representing a volume multicast relationship between a respective primary file system volume specified by a volume name (volumeName) and a respective group of secondary

1 data movers specified by an RCP group name (rcpgpeName). For example, at
2 configuration time, one or more RCP groups are defined in response to configuration
3 commands such as:

4
5 .RCP_config <server_name> add <IP>

6
7 This configuration command adds the IP address (IP) of a specified destination data
8 mover (server_name) to an RCP group.

9 Also at configuration time, a specified data mover can be defined to be a primary
10 data mover with respect to the RCP group (a relationship called a MultiCastNode) in
11 response to a configuration command such as:

12
13 .server_config <server_name> rep <groupname> add <IP>

14
15 where "server_name" is the name for the primary data mover, "groupname" is the name
16 of a configured RCP group, and "IP" is the IP address of the primary data mover. When
17 configuration of the MultiCastNode object is finished, the MultiCastNode object will
18 have its own name, a name for the primary data mover, an RCP group name, and a list of
19 IP addresses to which the primary server should broadcast in order to transmit IP packets
20 to all the secondary data movers in the RCP group.

21 The VolMCast object can then be built on top of a MultiCastNode object. The
22 additional information required for the VolMCast object is, on the sender side, the
23 primary or source file system volume and on each receiver side, the secondary or

1 destination file system volume. For flexibility, it is permitted to specify a different
2 volume name on each secondary data mover. By specifying the destination volume
3 names during creation of the VolMCast object, it is not necessary to specify the
4 destination volume names at each copy time. For example, the VolMCast object is
5 defined by configuration commands to the primary data mover such as:

```
6  
7 .server_config <server_name> "volmcast MultiCastNodeName"  
8 [ -src | -dest ] volume"
```

9
10 where <server_name> is the name of the MultiCast Node.

11 Once the VolMCast object has been defined, an IP-replication service can be
12 configured for the object upon the primary data mover. Then the primary data mover will
13 respond to commands for starting the replication service and stopping the replication
14 service upon the VolMCast object. When replication is stopped on a secondary, the
15 secondary file system is left in a consistent state. In other words, if a replay was in
16 progress, the stop will complete when the replay is finished.

17 The primary data mover may respond to additional commands for create a new
18 delta set on demand, updating the replication policy (high water mark and timeout
19 interval parameters) on the primary file system or secondary file systems, and defining
20 persistency of the replication process upon remount or reboot of the primary file system
21 or any one of the secondary file systems. For example, at reboot the replication service is
22 re-started on the primary file system and the secondary file system in the state it was at
23 unmount or shutdown. A recovery of the replication context happens at reboot or on

1 remount. The replica recovery is executed before the primary and secondary file systems
2 are made available for user access. This allows all modifications during the recovery of
3 the primary file system to be logged by the replication service.

4 As shown in FIG. 13, the volume multicast layer is responsive to a number of
5 commands 371 from higher layers in the protocol stack. In addition to the configuration
6 commands for defining a new VolMCast object relating a specified primary file system
7 volume to a specified RCP group, an existing VolMCast object can be opened for either a
8 sender mode or a receiver mode. An opened VolMCast object can be closed. Once a
9 VolMCast object has been opened in a sender mode, it can be called upon to broadcast a
10 control block (CB) to the secondary volumes of the VolMCast object, such as a control
11 block specifying a remote copy of a specified extent of the primary volume.

12 Control blocks may specify various operations upon the secondary volumes of the
13 VolMCast object, such as cluster file system commands for performing operations such
14 as invalidations, deletions, renaming, or other changes in the configuration of the objects
15 of the file system upon all copies (local or remote) of the file system. In this case, RCP is
16 used for the broadcast or forwarding of the cluster file system commands to all the data
17 movers that are to operate upon the local or remote copies of the file system, and for
18 returning acknowledgement of completion of the operations upon all of the copies of the
19 file system.

20 With reference to FIG. 13, the volume multicast layer defines a VolMCastSender
21 object 372 instantiated when a VolMCast instance is opened in the sending mode, and a
22 VolMCastReceiver object 373 instantiated when a VolMCast instance is opened in a
23 receiving mode. The VolMCastSender object class and the VolMCastReceiver object

1 class inherit properties of the VolMCast object class. When the volume multicast layer is
2 called upon in a primary data mover to maintain remote copies of a specified extent of a
3 VolMCastSender instance, an instance of a VolMCastCopy thread 374 is created and
4 executed. The VolMCastCopy thread instance accesses the delta sets from a primary
5 save volume 375 to produce a write stream 376 of blocks sent down to the RCP layer. At
6 the secondary data mover, an instance of a VolMCastReceiver thread 377 is instantiated
7 and executed to receive a read stream 378 of blocks and write the copied delta sets into a
8 secondary save volume 379. An instance of an acknowledgement thread 380 returns an
9 acknowledgement 381 of completion of copying of a delta-set for an extent to the
10 secondary file system. The acknowledgement is sent down to the RCP layer of the
11 secondary data mover. At the primary, the RCP layer sends the acknowledgement 382 to
12 an instance of an acknowledgement thread 383.

13 RCP is a session-layer protocol, for replication from one primary to multiple
14 secondary sites. Control is initiated by the primary, except when recovering from aborted
15 transfers. RCP uses TCP between the primary and secondary for control and data.
16 Network distribution is by an application-level multicast (ALM) using the RCP as a
17 forwarder. Port sharing with HTTP is used for crossing firewalls.

18 RCP may support other replication applications in addition to 1-to-N IP-based
19 replication for wide-area distribution of read-only data. These other applications include
20 1-to-N volume mirroring, cluster file system commands, remote file system replication,
21 and distribution and replication of other commands that may be recognized by the data
22 movers.

1 The 1-to-N volume mirroring is a simplification of to 1-to-N IP-based replication
2 for wide-area distribution of read-only data, because the volume mirroring need not
3 synchronize a remote volume with any consistent version of the primary volume until the
4 remote volume needs to be accessed for recovery purposes.

5 Remote file system replication also uses RCP for broadcast or forwarding an
6 application command to a remote data mover to initiate a replication of a file system
7 managed by the remote data mover. In a similar fashion, RCP may broadcast or forward
8 other commands recognized by data movers, such as iSCSI or remote-control type
9 commands for archival storage. For example, RCP could broadcast or forward remote
10 control commands of the kind described in Dunham, U.S. Patent 6,353,878 issued March
11 5, 2002 entitled "Remote Control of Backup Media in a Secondary Storage Subsystem
12 Through Access to a Primary Storage Subsystem," incorporated herein by reference.

13 The RCP forwarder is composed of two RCP sessions, an outbound session at the
14 primary, and an inbound session at the secondary. The inbound RCP session receives a
15 group name and looks up the group in a routing table. If routes for the group exist in the
16 routing table, then an RCP forwarder is created at the secondary, including a data path by
17 pointer passing from an "in" session to an "out" session.

18 An RCP group may be configured to include application-level multicast (ALM)
19 topology. For example, ALM route configuration commands begin with an identifier
20 number for the network file server ("cel") that contains the forwarder data mover, and an
21 identifier number ("ser") for the forwarder data mover in the network server. The
22 configuration commands end with a "nextthop" specification of an immediate destination
23 data mover:

1
2 cell1-ser2: rcproute add group=g1 nexthop=cell2-ser2

3 cell2-ser2: rcproute add group=g1 nexthop=cell2-ser3

4 cell2-ser2: rcproute add group=g1 nexthop=cell2-ser4

5
6 In effect, the forwarder data mover adds the "nexthop" specification to an entry
7 for the RCP group in the routing table in the forwarder data mover. This entire entry can
8 be displayed by the following configuration command:

9
10 cell2-ser2: rcproute display

11
12 The entry is displayed, for example, as a list of the "nexthop" destination data
13 movers. The entry can be deleted by the following configuration command:

14
15 cell2-ser2: rcproute delete

16
17 Each immediate destination data mover may itself be configured as a forwarder in
18 the RCP group. In this case, RCP commands and data will be forwarded more than once,
19 through a chain of forwarders. The set of possible RCP routes from a primary or
20 forwarder in effect becomes a tree or hierarchy of destinations.

21 The ALM commands may also include commands for creating sessions and
22 sending control blocks or data. For example, the following ALM command creates a
23 session and sends application data to all destinations in group "g1" from cell1-ser2 from a
24 file (named "filename") using a test application (named "rcptest").

cell-ser2: rcptest data=filename group=g1

FIG. 14 shows the RCP collector service 390 at a primary site. The programming for the RCP collector service includes an RCP session manager 391, collector and worker threads 392, and a single-thread RCP daemon 393. The RCP session manager 391 responds to requests from higher levels in the protocol stack, such as a request from an application 394 to open an RCP pipe 395 between the application 394 and the RCP collector service 390. The application 394 may then send to the session manager 391 requests to setup sessions with RCP groups. A session queue 396 stores the state of each session, and a control block queue 397 keeps track of control blocks sent via TCP/IP to the secondary data movers in the RCP groups. An RCP routing table 398 identifies the immediate destinations of each RCP group to which the TCP/IP messages from the RCP collection service are to be sent, as well as any other destinations to which the messages will be forwarded. For communication of the TCP/IP messages between the RCP service and the network, TCP port :80 is opened in both directions (i.e., for input and output). The single thread RCP daemon 393 is used for interfacing with this TCP port :80.

FIG. 15 shows the RCP collector service 400 at a secondary site. The RCP collector service at the secondary site is similar to the RCP collector service at the primary site, in that it includes an RCP session manager 401, collector and worker threads 402, a single thread RCP daemon 403 for access to/from TCP port :80, an RCP session state queue 406, an RCP control block queue 407, and an RCP routing table 408. The primary difference between the RCP collector service at the secondary site from the

1 RCP collector service at the primary site is in the collector and worker threads 402. At
2 the RCP secondary, the RCP commands and data are received from the TCP port :80
3 instead of from the application 404. The application 404 is the consumer of the RCP
4 data, instead of a source for RCP data. The RCP collector service 400 at the secondary
5 site may also serve as a forwarder for RCP commands, and therefore the RCP collector
6 service and worker threads 402 at the secondary site include a forwarder thread that does
7 not have a similar or complementary thread in the RCP collector service at the primary
8 site.

9 In operation, an application 404 can initialize the RCP collector service so that the
10 RCP collector service will call back the application upon receipt of certain RCP
11 commands from TCP port :80. For example, if a new connection command is received
12 from TCP port :80, then the RCP daemon 403 forwards the new connection command to
13 the RCP session manager. The RCP session manager 401 recognizes that this connection
14 command is associated with an application 404 at the secondary site, opens an RCP pipe
15 405 to this application, and calls the application 404 indicating that the RCP pipe 405 has
16 been opened for the RCP session. (The volume multicast receiver thread 377 of FIG. 13
17 is an example of such an application.) The application 404 returns an acknowledgement.
18 If the new connection is for a new RCP session, then the session manager creates a new
19 RCP session, and places state information for the new session on the RCP session queue
20 406. RCP control blocks and data may be received for the session from the TCP port :80.
21 The data may be forwarded to the application, or to a file specified by the application.
22 RCP control blocks to be executed by the RCP collector service 400 may be temporarily

1 placed on the control block queue 407. RCP control blocks or data intended for other
2 secondary site may be forwarded to the intended secondary sites.

3 FIG. 16 shows further details of the forwarding of RCP commands and data by a
4 data mover 430 identified as Cel2-Ser1. The data mover 430 is programmed with a
5 TCP/IP layer 431 for communication with the IP network 220, and an RCP layer 432
6 over the TCP/IP layer. For forwarding the RCP commands and data, the RCP layer 432
7 creates an inbound session 433 and an outbound session 434. The inbound session 433
8 receives RCP commands from the TCP/IP layer 431. The TCP/IP data stream is retained
9 in a data buffer 435. When an RCP command calls for the forwarding of RCP commands
10 or data to another data mover in a specified RCP group, the inbound session 433
11 performs a lookup for the group in a routing table 436.

12 In the example of FIG. 16, the routing table 436 includes a copy of all of the
13 routing information for each group of which the data mover 430 is a member. In this
14 case, for GROUP1, the primary data mover sends RCP commands to at least data movers
15 CEL2-SER1 and CEL9-SER1. CEL2-SER1 (i.e., the data mover 430) forwards the RCP
16 commands and RCP data to data movers CEL3-SER1 and CEL7-SER1. In particular, the
17 inbound session 433 creates an outbound session 434 and creates a TCP/IP data path
18 from the inbound session 433 to the outbound session 434 by passing pointers to the data
19 in the data buffer. The outbound session 434 invokes the RCP/IP layer 431 to multicast
20 the TCP data stream in the data buffer 435 over the IP network 220 to the data movers
21 CEL3-SER1 and CEL7-SER1.

22 The data mover CEL3-SER1 in succession forwards the RCP commands to data
23 movers CEL4-SER1 and CEL5-SER1. Normally, the data mover CEL2-SER1 (430)

1 does not need to know that the data mover CEL3-SER1 forwards the RCP commands to
2 data movers CEL4-SER1 and CEL5-SER1, but if the data mover CEL2-SER1 (430)
3 would fail to receive an acknowledgement from CEL3-SER1, then the data mover CEL2-
4 SER1 could minimize the impact of a failure of CEL3-SER1 by forwarding the RCP
5 commands to CEL4-SER1 and CEL5-SER1 until the failure of CEL3-SER1 could be
6 corrected.

7 FIG. 17 shows a flowchart of how the RCP collector service at the secondary site
8 processes an inbound RCP session command. In a first step 411, the RCP collector
9 service receives a session command. In step 412, if this session command is not a
10 command to be forwarded to other secondary sites, then execution branches to step 413
11 to execute the action of the command, and the processing of the session command is
12 finished.

13 In step 412, if the session command is a command to be forwarded to other
14 secondary sites, then execution continues from step 412 to step 414. In step 414, the
15 RCP collector service gets the RCP group name from the session command. Then, in
16 step 415, the RCP collector service looks up the group name in the RCP routing table
17 (408 in FIG. 15). If the group name is not found, then execution branches from step 416
18 to step 417. In step 417, the RCP collector service returns an error message to the sender
19 of the session command.

20 In step 416, if the group name is found in the RCP routing table, then execution
21 continues from step 416 to step 418. In step 418, the RCP collector service forwards the
22 action of the session command to each secondary in the group that is an immediate
23 destination of the forwarder (i.e., the data mover that is the secondary presently

1 processing the RCP session command). This is done by instantiating local replication
2 threads or creating outbound sessions for forwarding the action of the session command
3 to each secondary in the group that is an immediate destination of the forwarder. After
4 step 418, processing of the RCP session command is finished.

5 FIG. 18 shows an example of forwarding and local replication. In this example,
6 the IP network 220 connects a primary data mover 421 to a network file server 422 and a
7 secondary data mover 423. The network file server 422 includes three data movers 424,
8 425, and 426, and storage 427. The primary data mover manages network access to a
9 primary file system 428. The data mover 424 functions as a forwarder data mover. The
10 data mover 425 functions as a secondary data mover managing access from the network
11 to a secondary file system (copy A) 429. The data mover 426 functions as a secondary
12 data mover managing access from the network to a secondary file system (copy B) 430.
13 The data mover 423 manages network access to a secondary file system (copy C) 431.

14 In operation, when the primary data mover 421 updates the primary file system
15 428, it multicasts the modified logical blocks of the file system volume over the IP
16 network 220 to the forwarder data mover 424 and to the secondary data mover 423. The
17 forwarder data mover 424 receives the modified blocks, and performs a local replication
18 of the blocks to cause the secondary data mover 425 to update the secondary file system
19 (copy A) 429 and the to cause the secondary data mover 426 to update the secondary file
20 system (copy B) 430.

21 To perform the local replication, the forwarder data mover 424 has its volume
22 multicast layer (323 in FIG. 9) save the modified blocks in a save volume 432 in the
23 storage 427, and then the forwarder data mover 424 sends replication commands to the

1 local secondary data movers 425 and 426. Each local secondary data mover 425, 426 has
2 its playback module (328 in FIG. 9) replay the modifications from the save volume 432
3 into its respective secondary file system copy 429, 430.

4 FIG. 19 shows the sharing of the data mover's network TCP port :80 (440)
5 between HTTP and RCP. This configuration is used in all data movers having the RCP
6 collector service; i.e., primary, secondary, or forwarder. The TCP data channel from
7 TCP port :80 (440) provides an in-order byte stream interface. IP packets 444 for HTTP
8 connections and IP packets 445 for RCP connections from the network 220 are directed
9 to the data mover's TCP port :80 (440). The TCP port :80 (440) is opened in both
10 directions (i.e., input and output). In the input direction, the data mover uses a level 5
11 (L5) filter 441 for demultiplexing the IP packets for the HTTP connections from the IP
12 packets for the RCP connections based on an initial segment of each TCP connection.
13 The L5 filter hands the TCP connection off to either a HTTP collector service 442 or an
14 RCP collector service 443. (The RCP collector service 443 is the collector service 390 in
15 the RCP primary of FIG. 14 or the RCP collector service 400 in an RCP secondary of
16 FIG. 15.) For example, if the initial segment of a TCP connection contains "HTTP/1.X",
17 then the L5 filter 441 directs the IP packets for the connection to the HTTP collector
18 service 442. If the initial segment of the TCP connection contains "RCP/1.0", then the IP
19 packets for the TCP connection are directed to the RCP collector service 443. (In an
20 alternative arrangement, the connection could be split as is done in a conventional stand-
21 alone IP switch.)

22

Data Recovery With Internet Protocol Replication With Or Without Full

Resync

A remote replication system may protect against data loss resulting from a disaster. For example, FIG. 20 shows the state of a remote replication system before a disaster. The remote replication system copies data in an asynchronous fashion from a primary data storage system at the primary site to a secondary data storage system at a secondary site. The primary and secondary data storage systems could be file servers, for example as shown in FIG. 1.

As shown in FIG. 20, the primary site includes a replication service 450 replicating changes made to a primary file system 451, and the secondary site includes a playback service 452 writing the changes into a secondary file system 453. The replication service 450 transfers the changes in "delta sets" sent over an IP pipe 454 to the playback service 452. The disaster occurs at the primary site during this transfer. At this time, the playback service 452 is a few delta sets behind the replication service 450; for example, the playback service is playing back delta set (n-x) when the replication service 452 is replicating delta set (n).

Once the disaster causes the primary site to go down, the secondary site can be activated to service client read/write requests. When the primary site comes back up, it can be re-synchronized to the secondary site, in many cases without making a full copy of the file system. Then replication can be re-established as it was before the disaster. This recovery process has been designed to minimize data loss and the down time of both the primary and secondary site.

1 FIG. 21 is a flowchart of the preferred recovery process. Initially, in step 461,
2 data is replicated from the primary file system at the primary site by sending delta sets to
3 the secondary file system at the secondary site. Then in step 462, the primary site
4 becomes inoperative. In response, in step 463, the secondary site is activated for
5 read/write access to the secondary file system. This is done by a subroutine for failover
6 with a checkpoint and without sync, as further described below with respect to FIG. 23.
7 Eventually, in step 464, the primary site becomes operative. In response, in step 465, the
8 primary file system is synchronized to the state of the secondary file system, as further
9 described below with respect to FIG. 28. Then in step 466, read/write access to the
10 primary file system and replication of data from the primary file system to the secondary
11 file system is resumed in a failback operation, as further described below with reference
12 to FIG. 30.

13 FIG. 22 shows the state of the remote replication system of FIG. 20 during the
14 failover operation. The (n-x) deltasets are played back into the secondary file system
15 453, to put the secondary file system 453 into a state called a "restart point". A snapshot
16 copy facility 456 creates a snapshot 455 of the "restart point" in order to facilitate
17 synchronization of the primary file system with the secondary file system once the
18 primary site becomes operative. The playback service 452 is stopped, and the secondary
19 file system 453 is remounted as read/write. The snapshot copy facility 456 keeps a
20 record of changes made to the secondary file system 453 after the "restart point."

21 FIG. 23 shows a flowchart of the failover operation. In step 471, if possible, the
22 primary file system is remounted as a read-only file system. Next, in step 472, the
23 secondary site finishes playback of the (n-x) delta sets that it had received but had not

1 played back prior to the primary site becoming inoperative. In step 473, the secondary
2 site creates a snapshot copy of the "restart point" of the secondary file system. This
3 snapshot is stamped with the last delta set having been replayed into the secondary file
4 system prior to the creation of the snapshot. In other words, an identifier of this last delta
5 set is stored as an attribute of the snapshot. In general, whenever a snapshot is made of a
6 file system during playback of delta sets into the file system, the snapshot is stamped with
7 the last delta set having been replayed. In step 474, the secondary file system is
8 remounted as a read/write file system. In step 475, the snapshot process retains and
9 identifies all changes (delta) made to the secondary file system since the restarting point.

10 Preferably the snapshot copy facility 456 retains and identifies the changes at a
11 logical volume level of data storage. For example, the present state of the secondary file
12 system is stored in a "clone volume," and old versions of the logical blocks that have
13 been changed in the clone volume are saved in a "save volume". In order to conserve
14 storage, the logical blocks of the save volume are dynamically allocated to the old
15 versions of the changed blocks as the changes are made to the clone volume. As shown
16 in FIG. 24, for each logical block that has been changed in the clone volume, a block map
17 480 identifies the logical block address (S_i) of the old version of the block in the save
18 volume and the corresponding logical block address (B_i) of the changed block in the
19 clone volume.

20 FIG. 25 shows details of the preferred snapshot copy facility 456, which provides
21 multiple snapshots 483, 503 of a production file system 481. (Two successive snapshots
22 of the secondary file system are used, for example, in the resync subroutine further
23 described below with reference to FIG. 28.) The content of each snapshot file system

1 483, 503 is the state of the production file system 481 at a particular point in time when
2 the snapshot was created. The snapshot copy facility 456 provides a hierarchy of objects
3 in a volume layer 490 supporting the file systems in a file system layer 491. The
4 production file system 481 is supported by read/write access to a file system volume 482.
5 Each snapshot file system 483, 503 provides read-only access to a respective snapshot
6 volume 484, 504.

7 Additional objects in the volume layer 490 of FIG. 25 permit the content of each
8 snapshot file system to be maintained during concurrent read/write access to the
9 production file system 481. The file system volume 482 is supported by a snapped
10 volume 485 having read access to a clone volume 487 and write access to a delta volume
11 486. The delta volume 486 has read/write access to the clone volume 487 and read/write
12 access to a save volume 488.

13 In the organization of FIG. 25, the actual data is stored in blocks in the clone
14 volume 487 and a respective save volume 488, 506 in storage for each snapshot. The
15 delta volume 486 also accesses information stored in a bit map 489 and the block map
16 480. The bit map 489 indicates which blocks in the clone volume 487 have prior
17 versions in the save volume 488. In other words, for read-only access to the snapshot file
18 system, the bit map 489 indicates whether the delta volume should read each block from
19 the clone volume 487 or from the save volume 488. For example, the bit map is stored in
20 memory and it includes a bit for each block in the clone volume 487. The bit is clear to
21 indicate that there is no prior version of the block in the save volume 488, and the bit is
22 set to indicate that there is a prior version of the block in the save volume 488.

1 Consider, for example, a production file system 481 having blocks *a*, *b*, *c*, *d*, *e*, *f*,
2 *g*, and *h*. Suppose that when the snapshot file system 483 is created, the blocks have
3 values *a0*, *b0*, *c0*, *d0*, *e0*, *f0*, *g0*, and *h0*. Thereafter, read/write access to the production
4 file system 481 modifies the contents of blocks *a* and *b*, by writing new values *a1* and *a2*
5 into them. At this point, the following contents are seen in the clone volume 487 and in
6 the save volume 488:

7
8 Clone Volume: *a1*, *b1*, *c0*, *d0*, *e0*, *f0*, *g0*, *h0*

9
10 Save Volume: *a0*, *b0*
11

12 From the contents of the clone volume 487 and the save volume 488, it is possible
13 to construct the contents of the snapshot file system 483. When reading a block from the
14 snapshot file system 483, the block is read from the save volume 488 if found there, else
15 it is read from the clone volume 487.

16 FIG. 25 further shows that a snapshot queue 500 maintains respective objects
17 supporting multiple snapshot file systems 483, 503 created at different respective points
18 in time from the production file system 481. In particular, the snapshot queue 500
19 includes a queue entry (J+K) at the tail 501 of the queue, and a queue entry (J) at the head
20 502 of the queue. In this example, the snapshot file system 483, the snapshot volume
21 484, the delta volume 486, the save volume 488, the bit map 489, and the block map 480
22 are all located in the queue entry at the tail 501 of the queue. The queue entry at the head
23 of the queue 502 includes similar objects; namely, the snapshot file system (J) 503, a

1 snapshot volume 504, a delta volume 505, a save volume 506, a bit map 507, and a block
2 map 508.

3 The snapshot copy facility 456 may respond to a request for another snapshot of
4 the production file system 481 by allocating the objects for a new queue entry, and
5 inserting the new queue entry at the tail of the queue, and linking it to the snapped
6 volume 485 and the clone volume 487. In this fashion, the save volumes 488, 506 in the
7 snapshot queue 500 are maintained in a chronological order of the respective points in
8 time when the snapshot file systems were created. The save volume 506 supporting the
9 oldest snapshot file system 503 resides at the head 502 of the queue, and the save volume
10 488 supporting the youngest snapshot file system 483 resides at the tail 501 of the queue.

11 FIG. 26 shows a procedure performed by the snapshot copy facility for writing a
12 specified block (B_i) to the production file system. In step 511, if the snapshot queue is
13 not empty, execution continues to step 512. In step 512, the bit map at the tail of the
14 snapshot queue is accessed in order to test the bit for the specified block (B_i). Then in
15 step 513, if the bit is not set, execution branches to step 514. In step 514, the content of
16 the specified block (B_i) is copied from the clone volume to the next free block in the save
17 volume at the tail of the snapshot queue. Execution continues from step 514 to step 515.
18 In step 515, the save volume block address (S_i) of the free block is inserted into the entry
19 for the block (B_i) in the block map at the tail of the queue, and then the bit for the block
20 (B_i) is set in the bit map at the tail of the queue. After step 515, execution continues to
21 step 516. Execution also continues to step 516 from step 513 if the tested bit is found to
22 be set. Moreover, execution continues to step 516 from step 511 if the snapshot queue is

1 empty. In step 516, new data is written to the specified block (B_i) in the clone volume,
2 and then execution returns.

3 FIG. 27 shows a procedure performed by the snapshot copy facility for reading a
4 specified block (B_i) from a specified snapshot file system (N). In the first step 521, the
5 bit map is accessed for the queue entry (N) to test the bit for the specified block (B_i).
6 Then in step 522, if the tested bit is set, execution continues to step 523. In step 523, the
7 block map is accessed to get the save volume block address (S_i) for the specified block
8 (B_i). Then in step 524 the data is read from the block address (S_i) in the save volume,
9 and then execution returns.

10 If in step 522 the tested bit is not set, then execution branches to step 525. In step
11 525, if the specified snapshot (N) is not at the tail of the snapshot queue, then execution
12 continues to step 526 to perform a recursive subroutine call upon the subroutine in FIG.
13 27 for read-only access to the snapshot ($N+1$). After step 526, execution returns.

14 If in step 525 the snapshot (N) is at the tail of the snapshot queue, then execution
15 branches to step 527. In step 527, the data is read from the specified block (B_i) in the
16 clone volume, and execution returns.

17 FIG. 28 shows the subroutine to resync the primary file system with the secondary
18 file system, as used in the flowchart of FIG. 21. In a first step 531 in FIG. 28, the delta
19 set identifier ($n-x$) for the restart point is read from the delta set attribute of the restart
20 point snapshot, and it is incremented by one to compute ($n-x+1$). The save volume at the
21 primary site is searched for the delta set ($n-x+1$). Execution continues from step 531 to
22 step 532 if the delta set $n-x+1$ is found in the save volume at the primary site. In this

1 case, the primary site should also have all of the delta sets from delta set $n-x+1$ to delta
2 set n , and the primary file system can be restored to the state of the restart point in step
3 532 by an "undo" of the data blocks of these delta sets. The primary file system is
4 restored with "before images" of these data blocks in order to "undo" the changes. The
5 delta sets $n-x+1$ to n contain "after images" of these data blocks. The "before images" of
6 these data blocks are stored in the secondary site in the snapshot of the secondary file
7 system at the restart point. Therefore, to restore the primary file system to the state of the
8 restart point, the snapshot copy facility at the secondary site reads the "before images"
9 from the "restart point" snapshot, and the secondary site executes a remote copy of these
10 "before images" to the primary file system.

11 In step 532, the primary site restores the primary file system to the state of the
12 restarting point by obtaining a list of blocks from the save volume at the primary site,
13 including the blocks in delta set $n-x+1$ to delta set n . The primary site sends this list to
14 the snapshot copy facility at the secondary site. The secondary site retrieves the data of
15 these blocks from the snapshot at the restart point, and returns the data to the primary site.
16 The primary site receives these blocks of data and restores them to the primary file
17 system.

18 In step 533, the snapshot copy facility starts replication to the primary file system,
19 and creates an intermediate point snapshot (at state delta set $n-x+\delta$), so all new changes
20 made to the secondary file system since the intermediate point are kept and sent to the
21 primary file system. However, they are not yet played back into the primary file system.
22 At this point, the primary file system is mounted as a "raw" file system, so it is not
23 accessible to users.

1 In step 534, the changes made to the secondary file system from the restart point
2 to the intermediate point are copied from the secondary file system to the primary file
3 system. These changes are maintained by the snapshot copy facility at the secondary site.
4 For the snapshot copy facility as shown in FIG. 25, for example, assume that the restart
5 point is the snapshot J and the intermediate point is the snapshot J+K where K is one. In
6 this case, the changed blocks are identified in the bit map 507, and the data of the
7 changed data blocks are found by reading from the intermediate point snapshot file
8 system. To retrieve these changes, the snapshot copy facility scans the bit map 507 to
9 find any set bit indicating a change in the save volume, and if the bit is set, the snapshot
10 copy facility notes the block address (B_i). The snapshot copy facility then reads the
11 snapshot file system (J+K) 483 at this block address (B_i). The snapshot copy facility
12 continues this process until it has finished scanning the bit map 507.

13 In FIG. 28, after step 534, execution continues to step 535. In step 535, the
14 primary file system is mounted as a read-only UxFS file system, and playback to the
15 primary file system is started. At this point, the resync is finished, and the primary and
16 secondary sites are ready for failback.

17 In step 531, if the delta set $n-x+1$ cannot be found in the save volume from before
18 failover (459 in FIG. 29) at the primary site (for example, because the save volume was
19 destroyed during the disaster), then execution branches from step 531 to step 536. In this
20 case, the secondary file system is migrated to the primary file system, and a warning
21 message is logged and displayed to the system administrator to signal that the full copy
22 was required. In step 536, the secondary creates a first intermediate point snapshot (at
23 state delta set $n-x+\delta$), and copies the first intermediate point snapshot to the primary file

1 system. At this time, read/write access to the secondary file system is permitted on a
2 priority basis. Then in step 537, the secondary starts replication to the primary file
3 system, and creates a second intermediate point snapshot, so all new changes made to the
4 secondary file system since the second intermediate point are kept and sent to the primary
5 file system by the replication process. However, these new changes are not yet played
6 back into the primary file system. At this point, the primary file system is mounted as a
7 "raw" file system, so it is not accessible to users. Moreover, in step 537, the differential
8 of the first and second intermediate point snapshots (i.e., the changes made to the
9 secondary file system from the first intermediate point snapshot to the second
10 intermediate point snapshot) are copied to the primary file system. Then execution
11 continues to step 535, to mount the primary file system as a read-only UxFS file system,
12 and to start playback to the primary file system. Also in step 535, the restart snapshot
13 and the intermediate snapshot can be deleted.

14 The migration method of steps 536 to 537 is used in order to reduce the likelihood
15 that read/write access to the secondary file system and replication to the primary file
16 system would be interrupted by the primary and secondary save volumes (274, 275 in
17 FIG. 5) becoming full. Each save volume functions as a wrap-around buffer. The
18 primary save volume will store updates received in step 537 from the secondary site until
19 playback to the primary file system is started in step 535. If the primary save volume
20 would become full in step 537, a TCP flow control signal is sent from the primary site to
21 the secondary site to suspend transmission of further updates. When transmission of
22 further updates is suspended, the further updates can be stored in the secondary save
23 volume (275 in FIG. 5) until the secondary save volume would become full. If the

1 secondary save volume would become full, then read/write access to the secondary file
2 system and replication to the primary file system would be suspended until playback to
3 the primary file system is started in step 535 and updates are removed from the primary
4 and secondary save volumes.

5 FIG. 29 shows the state of the remote replication system of FIG. 20 during the
6 resync procedure. At the secondary site, the snapshot copy facility 456 has created the
7 intermediate point snapshot 457. A list of blocks in the delta sets n to $n-x+1$ is obtained
8 from the save volume 459 at the primary site and sent to the snapshot copy facility 456 at
9 the secondary site. The snapshot copy facility 456 returns "before images" of the
10 requested blocks over an IP pipe 458 to the primary file system 451 to restore the primary
11 file system to the state of the restart point snapshot, and then sends the changes from the
12 restart point to the intermediate point. The replication service 452 at the secondary site is
13 replicating changes to the secondary file system since the intermediate point, and is
14 sending these changes over the IP pipe 460 to the playback service 450. Transmission of
15 these changes since the intermediate point over the IP pipe 460 is concurrent with
16 transmission of changes prior to the intermediate point over the IP pipe 458. Once the
17 primary file system 451 has been synchronized to the intermediate point 457, the
18 playback service 450 at the primary site is activated to play these changes into the
19 primary file system 451. If the primary site should crash during the copy process of steps
20 534 or 536, then the copy process can be restarted automatically.

21 FIG. 30 is a flowchart of the subroutine to failback to the primary file system. In
22 a first step 541, the primary file system is made almost identical to the secondary file
23 system by the replication process. For example, when the primary file system is within

1 one delta chunk of the secondary file system, execution continues to step 542. In step
2 542, the secondary file system is re-mounted as read-only, and the last delta chunk is
3 copied from the secondary file system to the primary file system in order to synchronize
4 the primary file system from the secondary file system. Once synchronized, the primary
5 and secondary are aborted, in order to terminate the processes of replication from the
6 secondary, playback to the primary, and copying from the secondary to the primary. At
7 this point, the primary file system contains all the data that the secondary file system had
8 before the disaster and the additional data that the secondary file system had accumulated
9 after the disaster. Then in step 543, there is a resumption of the replication of data from
10 the primary file system and playback to the secondary file system. Then in step 544, the
11 primary file system is remounted as read/write. Therefore, the recovery process permits
12 replication to be restarted as it was before the disaster. Some data might have been lost at
13 failover time since replication is asynchronous. This is dependent on the number "x" of
14 delta sets that the secondary file system was behind the primary file system when the
15 disaster occurred at the primary. The less the number "x" of delta sets the closer the final
16 recovered file system would be to the original file system.

17 A number of commands have been devised for conducting the above recovery
18 procedure when network file servers (such as shown in FIGS. 1 and 5) are used at the
19 primary and secondary sites. These commands include a \$fs_replicate-failover
20 command, a \$fs_replicate -resync command, a \$fs_replicate-status command,
21 \$fs_replicate-failback command, and a \$fs_copy -start command.

22
23 The \$fs_replicate-failover command has the following format:

1
2 `$fs_replicate -failover <pfs>:cel=<server> <sfs> [-sync] [-now]`

3
4 where `<pfs>` is the primary file system, `<server>` is the server of the primary file system,
5 `<sfs>` is the secondary file system, `[-sync]` is an option that is set when the user wants the
6 failover to happen when both the primary file system and the secondary file system are
7 "in sync" (i.e., identical state), and `[-now]` is an option to specify immediate failover.
8 The `$fs_replicate-failover` command is executed on the control station for the secondary
9 file system.

10 When the `[-sync]` option is not set, this command will internally stop the
11 replication and playback of `<sfs>`, making sure that all of the delta sets available on
12 secondary site are re-played. Then it will try to stop the replication on `<pfs>` as well,
13 unless the `[-restartpoint]` option is set. Then `<sfs>` will be remounted "rw". If possible,
14 `<pfs>` is remounted "ro". No `[-sync]` option should be specified if the primary control
15 station or server is down. In this case, after failover, some data could be lost.

16 When the `[now]` option is set, the playback of `<sfs>` is immediately stopped,
17 without replaying any of the delta sets available on the secondary site.

18 When the `[-sync]` option is specified, the primary control station and the server
19 must be up, else the command will fail. The primary file system `<pfs>` is re-mounted as
20 read-only and a last delta set is created. Once the last delta set is played back to the
21 secondary, the replication process is stopped on the `<sfs>/<pfs>`. The secondary file
22 system `<sfs>` now contains all the data that the `<pfs>` had before the failover. The
23 secondary file system `<sfs>` is now mounted "rw".

1 FIG. 31 shows a flowchart of execution of the \$fs_replicate-failover command.
2 In a first step 551, if the sync option is not set, then execution branches to step 552 to
3 perform a failover without sync procedure, as described below with reference to FIG. 32.
4 After step 552, execution is finished. In step 551, if the sync option is set, then execution
5 continues to step 553. In step 553, if the primary site is not operational, execution returns
6 reporting a command failure. Otherwise, if the primary site is operational, execution
7 continues to step 554, to perform a failover with sync procedure, as described below with
8 reference to FIG. 33.

9 FIG. 32 shows a flowchart of the failover without sync procedure. In step 561, if
10 possible, the primary file system is remounted as a read-only file system. Then in step
11 562, the secondary site finishes playback of the delta sets that it had received but had not
12 played back prior to receipt of the failover command. In step 563, the secondary site
13 creates a snapshot copy of the secondary file system. This snapshot is called the
14 restarting point of the file system. In step 564, the secondary file system is remounted as
15 a read/write file system, and execution of the failover command is finished.

16 FIG. 33 shows a flowchart of the failover with sync procedure. In a first step 571,
17 the secondary file system is made almost identical to the primary file system by the
18 replication process. For example, once the secondary file system is within one delta set
19 chunk of the primary file system, execution continues to step 572. In step 572, the
20 primary file system is remounted as read-only, the last delta chunk is copied from the
21 primary file system to the secondary file system in order to synchronize the secondary
22 file system to the primary file system, and then the primary and secondary are aborted to
23 terminate the processes of replication from the primary, playback to the secondary, and

1 copying from the primary to the secondary. Then in step 573, the secondary site creates a
2 snapshot copy of the secondary file system. This snapshot copy is called the restarting
3 point of the file system. Finally, in step 574, the secondary file system is remounted as
4 read/write.

5 The `$fs_replicate-resync` command has the following format:

6
7 `$fs_replicate -resync sfs pfs:cel=server1`

8
9 where `sfs` is the secondary file system (this is now the source file system), `pfs` is the
10 primary file system (this is the file system to be restored, a raw file system mounted on a
11 server which has its IP interfaces configured), `server1` is the site where a disaster
12 occurred, and `<ckpt_name>` is the restart point name.

13
14 The `$fs_replicate-status` command has the following format:

15
16 `$fs_replicate -status <pfs> <sfs>:cel=<remote_server>`

17
18 where `<pfs>` is the primary file system, `<sfs>` is the secondary file system, and
19 `<remote_server>` is the server of the secondary site. This command can be used on the
20 control station for the primary site, and also on the control station for the secondary site.
21 Typically this command is used on the primary site when the user wants to check the
22 status of the replication process, to determine when it is reasonable to failback. It is

considered reasonable to failback when the difference between the primary file system and the secondary file system is very small. To assist the user, the primary site responds to this command by providing an indication of the number of blocks that are different between the primary file system and the secondary file system.

Following is an example of the information provided when the \$fs_replicate-status command is executed:

```
[nasadmin@cel_sfs nasadmin]$ fs_replicate -info sfs10
```

```
-v 128
```

```
id = 59
```

```
name = sfs10
```

```
type = replication
```

```
current_delta_set = 16
```

```
current_block_number = 1088
```

```
current_transfer_rate = 17.0667 blocks/second
```

```
avg_transfer_rate = 25.7444 blocks/second
```

```
id = 31
```

```
name = pfs10:cel_ip2
```

```
type = playback
```

```
current_delta_set = 16
```

```
outstanding delta sets:
```

```
Delta Source_create_time Blocks
```

```
-----
```

```
15 11/26 20:11:56 1
```

		Source			Destination				
	Delta	Create	Dur	Blocks	Playback	Dur	Blocks	Dsin	
		Time			Time			Group	
	----	-----			-----				
5	15	11/26 20:11:56	0	1					
6	14	11/26 20:01:56	0	1					
7	13	11/26 19:51:56	0	1	11/26 21:08:48	0	1	2	
8	12	11/26 19:41:56	0	1					
9	11	11/26 19:31:56	0	1	11/26 20:48:48	0	1	2	
10	10	11/26 19:21:56	0	1					
11	9	11/26 19:11:56	0	1	1 /26 20:28:28	0	1	2	
12	8	11/26 18:52:40	0	1	11/26 20:08:09	0	1	1	
13	7	11/26 18:42:40	0	1					
14	6	11/26 18:32:40	0	1	11/26 19:48:09	0	1	2	
15	5	11/26 18:22:39	0	1					
16	4	11/26 18:12:39	0	1	11/26 19:28:08	0	1	2	
17	3	11/26 18:02:39	0	4	11/26 19:07:23	0	4	1	

18

19 The \$fs_replicate-failback command has the following format:

20

21 \$fs_replicate -failback <pfs> <sfs>:cel=<remote_server>

22

23 where <pfs> is the primary file system (a UxFS file system mounted as "ro"), and <sfs>

24 is a secondary file system. This command is used when the user wants to failback to the

primary site (after re-build phase). No data is lost and the replication will continue but it will be reversed (the primary file system, which runs playback, will then run replication and vice-versa for the secondary file system). This command can also be used to change the direction of replication, in the absence of any failure or disruption.

The `$fs_replicate--failback` command is executed on the site that the replication service is running (i.e., the site having the file system that is mounted as read/write). Typically, the `$fs_replicate--failback` command is executed on the primary site after it has been rebuilt. Then the primary file system is a raw file system that is restored from the secondary file system. The secondary file system is re-mounted read write. Replication is stopped on the secondary and playback is stopped on the primary. After the primary file system and the secondary file systems are in sync, the primary file system is converted to a UxFS file system and re-mounted as read/write. Replication is turned back on and the flow of data is now from the primary file system to the secondary file system. Playback is started on the secondary file system. This command can be used to "swap" primary and secondary without having to do a full copy and without having to restart replication. On error, it should be possible to re-issue the `fs_replicate -failback` command and proceed if the error condition has been cleared.

The `$fs_copy--start` command has the following format:

```
$fs_copy --start <fs_ckpt2> <sfs>:cel=<remote_server> -fromfs <fs_ckpt1>
```

where `<fs_ckpt1>` is the primary file system ckpt1, `<fs_ckpt2>` is the primary file system ckpt2, `<sfs>` is the secondary file system, and `<remote_server>` is the server of the

1 secondary site. This command is used internally by the fs_replicate-resync command. It
2 could also be used for other purposes, for example, as described below with reference to
3 FIGS. 36-38. This command copies the delta from fs_ckpt1 to fs_ckpt2 over to the
4 secondary file system, knowing that the state of the secondary file system was equal to
5 fs_ckpt1.

6 The following is an example of how the commands can be used to recover from a
7 disaster.

- 8 1. After the disaster, decide the option to use for failover.
- 9 2. `$fs_replicate -failover pfs:cel=cel_pfs sfs`
- 10 3. Verify that sfs can accept read/write operation.
- 11 4. Initialize replication from sfs to pfs:

12 On PFS control station:

13 `server_mount server_2 -o ro pfs /pfs`

14 On SFS control station:

15 `$fs_replicate -resync pfs:cel=cel_pfs sfs`

- 16 5. Verify that replication from sfs to pfs is running without a problem.
- 17 6. `$fs_replicate -failback sfs pfs:cel=cel_pfs`
- 18 7. Verify that replication from pfs to sfs is running without a problem.

19

20 **Replication of Snapshots Using IP File System Copy Differential**

21 As described above with reference to steps 534 and 537 of FIG. 28, a file system
22 copy differential was used to replicate an intermediate point snapshot during recovery of

1 a file system from a disaster. A general-purpose file system copy differential command
2 (\$fs_copy-start) was also introduced.

3 FIG. 34 shows a flowchart of a procedure for the \$fs_copy -start command. Here
4 L is the index of the snapshot fs_ckpt1, and M is the index of the snapshot fs_ckpt2. In a
5 first step 581, M is compared to L, and execution branches to step 582 to return an error
6 if M is not greater than L. Otherwise, execution continues to step 583.

7 The following steps execute the \$fs_copy -start command by scanning one or
8 more of the bit maps of the snapshot copy facility of FIG. 25. Each bit map indicates
9 whether or not each block of a respective snapshot changed from the snapshot to the next
10 snapshot. Therefore, by examining the bit maps for the snapshots L to M-1, it is
11 determined whether or not each block has changed from snapshot L to M. If so, the new
12 value of the block is read from the snapshot M. The bit maps are scanned in block order
13 so that the new blocks are read from the snapshot M of the primary (i.e., the source) file
14 system and transmitted to the secondary (i.e., the destination) file system in the order of
15 their block indices (B_i).

16 In step 583, the block index (B_i) is set to zero. In step 584, a snapshot index (I) is
17 set to L. In step 585, the bit map for the snapshot (I) is indexed with the block index (B_i)
18 to determine whether or not the block was changed between snapshot (I) and snapshot
19 (I+1). If in step 585 the bit for the block (B_i) is not set in the bit map for the snapshot (I),
20 then no such change occurred, and execution continues to step 586. In step 586, if the
21 snapshot index I is not equal to M-1, then execution continues step 587. In step 587, the
22 snapshot index I is incremented by one. After step 587, execution loops back to step 585.

23 In step 585, if the bit for the block (B_i) is set in the bit map for the snapshot (I),

1 then execution branches to step 589. In step 589, the snapshot (M) is read to get the new
2 data for the block (B_i). Execution continues to step 590 of FIG. 35. In step 590, the
3 block index (B_i) and the new data for the block (B_i) from the snapshot (M) are returned
4 for copying into the secondary file system. Then in step 591, if the block index (B_i) is at
5 the end of the production volume, then execution is finished. If not, execution loops back
6 to step 588 of FIG. 34. In step 588, the block index (B_i) is incremented by one, and
7 execution loops back to step 584, to continue scanning for the next value of the block
8 index.

9 In step 586 of FIG. 34, if I is equal to M-1, then execution continues to step 591
10 of FIG. 35.

11 It should be understood that the flowchart of FIGS. 34-35 represents a program
12 executed by at least one processor in a data storage system such as a network file server.
13 The processor, for example, is a data mover computer (e.g., 232 in FIG. 1). The program,
14 for example, is initially contained in a program storage device such as a floppy disk (e.g.,
15 238 in FIG. 1) and down-loaded into storage of the data mover computer.

16 The program in the flowchart of FIGS. 34-35 has an inner loop including steps
17 585, 586, 587 that indexes the snapshots L to snapshot M-1. This sequence includes the
18 snapshot L and the snapshots that are both younger than the snapshot L and older than the
19 snapshot M. The program in the flowchart of FIGS. 34-35 has an outer loop including
20 steps 584, 585, 586, 591, and 588 that indexes the blocks. When a bit in the indexed bit
21 map is found to be set in step 585, the inner loop is exited to return the block index (B_i)
22 and the data in the snapshot M for block (B_i).

1 The snapshot copy differential has been described above for facilitating recovery
2 of a file system after a disaster. The snapshot copy differential can also be used for wide-
3 area distribution of updates on an as-needed basis. This reduces network traffic for the
4 case where a client has an old local version of a file system and needs a new version of
5 the file system. A new local version of the file system can be constructed by copying the
6 appropriate changes into the old local version of the file system.

7 FIG. 36, for example, shows a block diagram of a data network in which snapshot
8 deltas are transmitted over a wide-area network 626 from a network file server 627 to a
9 local file server 624 in order to update a local file system 625 as needed. The local file
10 server 624 services local clients 621, 622, and 623. When a client needs a more recent
11 version of the file system, and the local file system 625 is not the most recent version,
12 then the local file server may request a specified version from the network file server 627,
13 or the client may request the most recent version available.

14 The network file server 627 has a snapshot copy facility 628 storing multiple
15 snapshots 629, 630. If the local file system 625 in the local file server 624 is one of the
16 multiple snapshots, then the network file server 627 may respond to a request from the
17 local file server 624 by obtaining from the snapshot copy facility 628 a snapshot copy
18 differential that would contain all of the updates necessary to convert the local file system
19 624 to a more recent snapshot of the file system. In the usual case, the local file server
20 624 would request all of the updates necessary to convert the local file system 625 to the
21 most recent snapshot copy. However, it is also possible for the local file server 624 to
22 request the updates for a specified version that would not necessarily be the most recent
23 snapshot copy.

FIGS. 37 and 38 show a flowchart of a procedure for the replication of the most recent snapshot in the system of FIG. 36 using the snapshot copy differential. In a first step 641, the client requests access to the file in the local file system in the local file server. In step 642, the local file server accesses attributes of the local file system and finds that it is version (Q) of a local updatable file system, last updated at time (Tu) from a network file server having a network identifier (NETID). The local file server uses the time of last update (Tu) to determine that it is time to check for a more recent version, and sends an update request to the network file server (NETID). For example, the file system has an attribute specifying a minimum update interval (Tmin), and it is not time to request an update unless the minimum update interval has elapsed since the time of the last update. The request specifies the version (Q) already in the local file server. In response, the network file server accesses the snapshot copy facility to find the oldest version (J) and the youngest version (J+K) stored in the network file server. In step 645, Q is compared to J+K. If $Q=J+K$, then execution branches to step 646 of FIG. 38. In step 646, the network file server returns a message that no more recent version is available. The local file server resets the time of last update (Tu) to the current time, and accesses the local version (Q) for the client. Upon completion of step 646, the snapshot replication process is finished.

If in step 645 Q is not equal to J+K, then execution continues to step 647 of FIG. 38. In step 647, Q is compared to K. If Q is less than K, then execution branches to step 648. In this case, the version (Q) is not a snapshot in the snapshot copy facility because the version (Q) is too old. In step 648, the network file server copies the youngest snapshot version (J+K) to the local file server. The local file server replaces the local

1 version (Q) with the new version (J+K), resets the time of last update (Tu) to the current
2 time, and accesses the new local version (J+K) for the client. Upon completion of step
3 648, the snapshot replication process is finished.

4 In step 647, if Q is not less than K, then execution continues to step 649. In step
5 649, the network file server does a file system copy snapshot delta <Q> to <J+K> of
6 blocks into the local version (Q) to convert it into the youngest snapshot version (J+K).
7 The local file server resets the time of last update (Tu) to the current time, and accesses
8 this local snapshot version (J+K) for the client. Upon completion of step 649, the
9 snapshot replication process is finished.

10 In a preferred snapshot copy facility, as described below with reference to FIGS.
11 41 to 46, there is kept a meta bit map for each snapshot copy for indicating blocks of the
12 production file system that are not used in the snapshot copy. Further details of such a
13 snapshot copy facility are described in Philippe Armangau, et al., "Data Storage System
14 Having Meta Bit Maps for Indicating Whether Data Blocks are Invalid in Snapshot
15 Copies," U.S. Patent Application Ser. 10/213,241 filed Aug. 6, 2002, incorporated herein
16 by reference. The snapshot copy facility maintains the meta bit maps in order to store the
17 "before image" of a block in the save volume at the tail of the snapshot queue only when
18 the block is being written to and the "before image" is needed for responding to any
19 request for reading a snapshot copy. This reduces the number of blocks that are stored in
20 the save volumes. However, in this case, the bit map for each snapshot (L) indicates
21 whether or not a block has been stored in the save volume for the snapshot (L), and no
22 longer will indicate all of the blocks that have been changed after snapshot (L) and before
23 snapshot (L+1). In particular, if a block was not in use for snapshot (L), and was written

1 to after snapshot (L) and before snapshot (L+1), then the "before image" of the block will
2 not be written into the save volume for snapshot (L).

3 When it is known that a block is not used in the snapshot copy (M), then there is
4 no need for the snapshot copy facility to return the block when responding to a request
5 for the snapshot delta of snapshot <L> to snapshot <M>. Therefore, for the preferred
6 snapshot copy facility, it is desirable to modify the procedure of FIG. 34 in order use the
7 information in the meta bit map for the snapshot <M>. In this case, the procedure of
8 FIG. 34 should also be modified to account for the fact that the save volumes no longer
9 store the "before images" for all of the blocks that may have changed between the
10 successive snapshot copies.

11 FIG. 39 shows how the flowchart of FIG. 34 can be modified for use with the
12 preferred snapshot copy facility of FIGS. 41 to 46. Steps 651, 652, and 653 of FIG. 39
13 are similar to steps 581, 582, and 583 of FIG. 34. Step 654 of FIG. 39 is similar to step
14 588 of FIG. 34, and step 658 of FIG. 39 is similar to step 589 of FIG. 34.

15 In FIG. 39, after step 653 or step 654, execution continues to step 655. In step
16 655, if the block (B_i) is not in use in the snapshot (M), then execution branches to step
17 591 of FIG. 35. Therefore, a block not in use in the snapshot (M) is not returned in
18 response to the command to copy the snapshot delta <L> to <M>. If the block (B_i) is in
19 use in the snapshot (M), then execution continues from step 655 to step 656. In step 656,
20 if the block (B_i) is in any of the save volumes (L) to (M-1), then execution continues to
21 step 658 to read the block (B_i) from the snapshot (M) in order to return this version of the
22 block (B_i) in response to the command to copy the snapshot delta <L> to <M>. In step
23 656, if the block (B_i) is in any of the save volumes (L) to (M-1), then execution branches

1 to step 657. In step 657, if the block (B_i) is in use in all of the snapshots (L) to (M-1),
2 then execution branches to step 591 of FIG. 35. In this case, the block (B_i) did not
3 change from snapshot (L) to snapshot (M), because if it did, a "before image" would have
4 been stored in one of the save volumes (L) to (M-1). In step 657 of FIG. 39, if it is not
5 true that the block (B_i) is in use in all of the snapshots (i.e., it is true that the block (B_i) is
6 not in use in at least one of the snapshots (L) to (M-1)), then execution continues from
7 step 657 to step 658 to read the block (B_i) from the snapshot (M) in order to return this
8 version of the block (B_i) in response to the command to copy the snapshot delta <L> to
9 <M>. In this case, it is possible that the block changed from snapshot (L) to snapshot
10 (M) despite the fact that the block (B_i) is not in any of the save volumes (L) to (M-1).

11 FIG. 40 shows a preferred implementation of the procedure of FIG. 39. Steps 661
12 to 665 of FIG. 40 are similar to steps 651 to step 655 of FIG. 39. In step 665, the meta bit
13 map for snapshot (M) has a value for the block (B_i) indicating whether or not the block
14 (B_i) is in use for the snapshot (M). In particular, a value of 1 indicates that the block (B_i)
15 is in use for the snapshot (M). Steps 666, 669, and 670 of FIG. 40 are similar to step 584,
16 586 and 587 of FIG. 34. From step 666 or step 670, execution continues to step 667.

17 In step 667, if the bit map for snapshot (I) has a value of 1 for the block (B_i), then
18 execution continues to step 671 to read the snapshot (M) to get data for the block (B_i), in
19 order to return the data in response to the command to copy the snapshot delta <L> to
20 <M>. In this case, the save volume for block (I) includes a "before image" for the block
21 (B_i). Otherwise, if the bit map for snapshot (I) does not have a value of 1 for the block
22 (B_i), execution branches from step 667 to step 668. In step 668, if the meta bit map for

1 the snapshot (I) does not have a value of 1 for the block (B_i), execution continues to step
2 671 to read the snapshot (M) to get data for the block (B_i), in order to return the data in
3 response to the command to copy the snapshot delta <L> to <M>. In this case, the block
4 (B_i) is not in use in the snapshot (I). Otherwise, if the meta bit map for the snapshot (I)
5 has a value of 1 for the block (B_i), execution continues to step 669.

6 7 Maintenance of Meta Bit Maps in the Snapshot Copy Facility

8 In the above description of the snapshot copy process, and in particular with
9 respect to FIG. 25, it was assumed that the original contents of a block of the production
10 file system must be saved to the most recent save volume before the contents of the block
11 are modified by a write access to the production file system. In practice, however, the
12 original contents are often invalid, and therefore need not be saved. For example, many
13 applications start with an empty file, and the file increases in size as data is written to the
14 file. In some of these applications, the file rarely decreases in size. However, storage for
15 the file may be released when the file is deleted from the file server, for example, when
16 the file is transferred to archival storage. In some applications, the extent of a file may be
17 dynamically decreased concurrent with read/write access to the file.

18 There are significant advantages to identifying when read/write access to the
19 production file system is about to modify the contents of an invalid data block. If this can
20 be done in an efficient manner, then there can be a decrease in the access time for write
21 access to the production file system. A write operation to an invalid block can be
22 executed immediately, without the delay of saving the original contents of the data block
23 to the most recent save volume at the tail of the snapshot queue. Moreover, there is a

1 saving of storage because less storage is used for the save volumes. There is also a
2 decrease in memory requirements and an increase in performance for the operations upon
3 the snapshot file systems, because smaller bit and block hash indices can be used, and the
4 reduced amount of storage for the snapshots can be more rapidly restored to the
5 production file system, or deallocated for re-use when snapshots are deleted.

6 An efficient way of identifying when read/write access to the production file
7 system is about to modify the contents of an invalid data block is to use a meta bit map
8 having a bit for indicating whether or not each allocated block of storage in the
9 production file system is valid or not. For example, whenever storage is allocated to the
10 production file system by the initial allocation or the extension of a clone volume, a
11 corresponding meta bit map is allocated or extended, and the bits in the meta bit map
12 corresponding to the newly allocated storage are initially reset.

13 FIG. 41 shows a procedure for writing a specified block (Bi) to the production file
14 system when there is a meta bit map for indicating invalid data blocks in the production
15 file system. In a first step 681, a queue pointer is set to point to the queue entry at the tail
16 of the snapshot queue. Next, in step 682, the bit map in this snapshot queue entry is
17 accessed to test the bit for the specified block (Bi). Next, in step 683, if the tested bit is
18 found to be set, then execution continues to step 684. In step 684, new data is written to
19 the specified block (Bi) in the clone volume, and then execution returns.

20 In step 683, if the bit in the bit map is not set, then execution branches to step 685.
21 In step 685, the meta bit map in the snapshot queue entry is accessed to test the bit for the
22 specified block (Bi). Then, in step 686, execution continues to step 687 if this bit is
23 found to be set. In step 687, the content of the block (Bi) is copied from the clone

1 volume to the next free block in the save volume at the tail of the snapshot queue. In step
2 688, an entry for the block (Bi) is inserted into the block map at the tail of the snapshot
3 queue, and then the bit for the block (Bi) is set in the bit map at the tail of the snapshot
4 queue. Execution continues from step 688 to step 684, to write new data to the block (Bi)
5 in the clone volume.

6 In step 686, if the tested bit is found not to be set, then execution branches to step
7 689. If the queue pointer is pointing to the head of the queue, then execution branches to
8 step 684, to write new data to the block (Bi) in the clone volume. Otherwise, if the queue
9 pointer is not pointing to the head of the snapshot queue, then execution continues to step
10 690 to advance the queue pointer to the next snapshot queue entry toward the head of the
11 snapshot queue. After step 690, execution loops back to step 682.

12 FIG. 42 shows an organization of the snapshots in the network file server when a
13 respective meta bit map 692, 693 is maintained for each snapshot in addition to the meta
14 bit map 691 for the production volume. It is desired to maintain a respective meta bit
15 map for each snapshot so that whenever the production file system is restored with a
16 snapshot file system, the meta bit map for the production file system can be restored with
17 the meta bit map for each snapshot. For example, a meta bit map of the production file
18 system is included in the production file system, so that whenever a snapshot copy of the
19 production file system is created, a snapshot copy of the meta bit map is also created.
20 Consequently, when the production file system is restored with a snapshot, the meta bit
21 map of the production volume is replaced with the meta bit map of the snapshot.

22 As shown in FIG. 42, a meta bit map 691 is linked to the production file system
23 481 for indicating invalid blocks in the production file system. Each entry in the

1 snapshot queue 500 includes a respective meta bit map linked to the snapshot file system
2 in the entry. For example, the queue entry (J+K) at the tail 501 of the queue has a meta
3 bit map 692 linked to the snapshot file system 483, and the queue entry (J) at the head
4 502 of the queue includes a meta bit map 693 linked to the delta volume 505.

5 To reduce the memory and storage requirements for maintaining the bit maps 694,
6 696 and block maps 695, 697, the each bit map is organized as a set of pages indexed by
7 a page table, and the each block map is organized as a set of hash lists indexed by a hash
8 table. The bit maps and block maps 694, 695 at the queue entry (J+K) at the tail of the
9 queue are initially created in a random access memory cache and written back to storage
10 of the save volume 488 when a next snapshot of the production file system 481 is created.
11 Thereafter the bit maps and block maps can be deallocated from the random access
12 memory in order to free up cache memory, and later staged from storage to the cache
13 memory when needed for read access to their respective snapshots.

14 FIG. 43 shows that the bit map 694 is organized as a page table 701 indexed by a
15 set of most significant bits (MSB) of the block index (B_i), and pages 702 and 703 of the
16 bit map linked to respective entries of the page table. The page table 701 includes a set
17 of entries, each of which is either zero, indicating that the entire page is zero and
18 therefore the page is not stored in storage or random access memory, or is a pointer to a
19 page of the bit map.

20 FIG. 44 shows that the block map 695 is organized as a hash table 708 indexed by
21 a hashing of the block index (B_i), and a set of hash lists 709 linked to respective entries
22 of the hash table 708. Each non-zero entry in the hash table 708 points to a respective
23 one of the hash lists 709. Each entry in each hash list includes a block address (B_i) to a

1 block in the clone volume, a corresponding block address (S_i) of the block in the save
2 volume, and a value that is either zero indicating the end of the has list, or a pointer to the
3 next entry in the list.

4 It is not necessary to retain the bit map 694 for the snapshot because the block
5 map 695 can be accessed to provide the same information that is contained in the bit map
6 694. In particular, the block map 695 can be accessed simply to determine whether or not
7 a specified block (B_i) is found in the block map, and therefore can be found in the save
8 volume for the corresponding snapshot. However, the bit map 694 can be accessed more
9 quickly to provide this information. Therefore, for applications such as backup where the
10 snapshots would be accessed very infrequently, it may be desirable to conserve storage
11 by discarding the bit map for the snapshot at the tail of the queue once the next snapshot
12 is created. Otherwise, for an application such as described above with respect to FIGS.
13 36-38 where the snapshots would be accessed frequently, the bit map for the snapshot at
14 the tail of the snapshot queue may be written to storage and maintained in a random
15 access cache memory as needed.

16 FIG. 45 shows that the meta bit map 691 has a respective bit corresponding to
17 each block in the clone volume, and in this example, each bit in the meta bit map
18 corresponds to one and only one block in the clone volume. The meta bit map 691
19 includes a series of words, each with a multiple of M bits. In this example, a bit having a
20 value of zero indicates a corresponding block that is invalid, and a bit having a value of
21 one indicates a corresponding block that is valid.

22 The meta bit map, however, may have a granularity greater than one block per bit.
23 For example, each bit in the meta bit map could indicate a range of block addresses,

1 which may include at least some valid data. The benefit to the increased granularity is a
2 reduced size of the meta bit map at the expense of sometimes saving invalid data to the
3 save volume. For example, FIG. 46 shows the interpretation of a meta bit map 691'
4 having a granularity of two blocks per bit. Each bit is set if any one of the two
5 corresponding blocks is valid, or conversely, each bit is clear only if neither of the two
6 corresponding blocks is valid. In this case, the block address can be converted to a bit
7 address by an integer division by two, for example, by an arithmetic right shift of the
8 block address by one bit position.

9
10 In view of the above, there has been described a snapshot copy facility that stores
11 a plurality of snapshot copies of a production file system. Each of the snapshot copies is
12 a prior state of the production file system at a respective point in time. The snapshot
13 copy facility receives a request for the difference between a specified older snapshot copy
14 and a specified younger snapshot copy, and responds by returning the difference between
15 the older snapshot copy and the younger snapshot copy. In a preferred implementation,
16 the snapshot copy facility has an index for each snapshot copy for indicating blocks of
17 data in the production file system that have changed between the snapshot copy and a
18 next snapshot copy of the production file system. The indices are scanned for a sequence
19 of the snapshot copies to determine the blocks that have changed, and the snapshot copy
20 facility returns the block numbers and data in the younger snapshot copy for the blocks
21 that have changed.

22 It has been shown how the difference between an older snapshot copy and a
23 younger snapshot copy can be used for recovery purposes and also for distribution of

1 updates over a network on an as-needed basis. In general, a client has an older local copy
2 of a snapshot and needs a more recent snapshot. The client requests a network file server
3 to provide the difference between the older snapshot and the more recent snapshot. The
4 client writes this difference into its older local copy to reconstruct the more recent
5 snapshot.